



(MURI-09) DISTRIBUTED LEARNING AND INFORMATION DYNAMICS IN NETWORKED AUTONOMOUS

Eric Feron
GEORGIA TECH RESEARCH CORPORATION

11/20/2015
Final Report

DISTRIBUTION A: Distribution approved for public release.

Air Force Research Laboratory
AF Office Of Scientific Research (AFOSR)/ RTA2
Arlington, Virginia 22203
Air Force Materiel Command

REPORT DOCUMENTATION PAGE					Form Approved OMB No. 0704-0188	
<p>The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to the Department of Defense, Executive Service Directorate (0704-0188). Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.</p> <p>PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ORGANIZATION.</p>						
1. REPORT DATE (DD-MM-YYYY) 09-11-2015		2. REPORT TYPE Final Report			3. DATES COVERED (From - To) From July 1, 2009 to June 30, 2015	
4. TITLE AND SUBTITLE DISTRIBUTED LEARNING AND INFORMATION DYNAMICS IN NETWORKED AUTONOMOUS SYSTEMS					5a. CONTRACT NUMBER	
					5b. GRANT NUMBER FA9550-09-1-0538	
					5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) Eyad Abed, Nina Balcan, John Baras, Eric Feron, Leslie Kaebing, Nuno Martins, Asuman Ozdaglar, Jeff Shamma, Emilio Frazzoli, Peyton Young					5d. PROJECT NUMBER	
					5e. TASK NUMBER	
					5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) Gerogia Tech Research Corporation 505 10th ST NW Atlanta GA 30318-5775					8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR/PKR3 CODE FA9550 6. USAF, AFRL DUNS 143574726 AF OFFICE OF SCIENTIFIC RESEARCH 875 NORTH RANDOLPH STREET, RM 3112 ARLINGTON VA 22203					10. SPONSOR/MONITOR'S ACRONYM(S)	
					11. SPONSOR/MONITOR'S REPORT NUMBER(S)	
12. DISTRIBUTION/AVAILABILITY STATEMENT Approved for Public Release						
13. SUPPLEMENTARY NOTES N/A						
14. ABSTRACT This project is motivated by the need to enable advanced operations of teams of autonomous vehicles to learn and adapt to uncertain and hostile environments under effective utilization of communications resources. Of particular interest is the interplay between distributed learning and information dynamics. Distributed learning refers to a collection of interacting agents with limited local processing, information, and communications, all seeking to achieve a global objective in an uncertain and possibly hostile environment. Information dynamics refers to the architecture, either inherited or designed, of information flow among the distributed agents. The interplay of distributed learning algorithms and information dynamics can have dramatic effects on the efficiency of the collective. Specific research thrusts include: i) online resource allocation, ii) networked operations, and iii) evolving and uncertain operations.						
15. SUBJECT TERMS Distributed Learning, Game theory, Information dynamics, uncertain operations.						
16. SECURITY CLASSIFICATION OF: a. REPORT U			17. LIMITATION OF ABSTRACT U		18. NUMBER OF PAGES 58	
b. ABSTRACT U					19a. NAME OF RESPONSIBLE PERSON Dr. Eric Feron	
c. THIS PAGE U					19b. TELEPHONE NUMBER (Include area code) 617 413 7666	

Reset

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATES COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33615-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report, e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5d. PROJECT NUMBER. Enter all project numbers as they appear in the report, e.g. 1F665702D1257; ILIR.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.

FINAL PERFORMANCE REPORT: DISTRIBUTED LEARNING AND INFORMATION DYNAMICS IN NETWORKED AUTONOMOUS SYSTEMS

AFOSR Grant #FA9550-09-1-0538

PI: Eric Feron (current)

Jeff S. Shamma (former)

Georgia Institute of Technology

Atlanta, GA 30332

1 Abstract

This project is motivated by the need to enable advanced operations of teams of autonomous vehicles to learn and adapt to uncertain and hostile environments under effective utilization of communications resources. Of particular interest is the interplay between distributed learning and information dynamics. Distributed learning refers to a collection of interacting agents with limited local processing, information, and communications, all seeking to achieve a global objective in an uncertain and possibly hostile environment. Information dynamics refers to the architecture, either inherited or designed, of information flow among the distributed agents. The interplay of distributed learning algorithms and information dynamics can have dramatic effects on the efficiency of the collective. Specific research thrusts include: i) online resource allocation, ii) networked operations, and iii) evolving and uncertain operations.

2 Status/Progress

We will present our research efforts on distributed learning and information dynamics according to the following interrelated research thrusts:

- *Distributed online resource allocation:* The general theme of distributed learning concerns how interacting decision making components can self-configure in the absence of a centralized planning authority. Of particular interest resource allocation problems, both because of their DoD relevance (e.g., weapon-target allocation, sensor coverage, distributed routing, etc.) and their underlying special structure. This thrust emphasizes distributed resource allocation through the use of game theoretic models and methods.
- *Networked operations:* The main implication of a distributed decision architecture is distributed information. Even in a multi-component setting (such as mobile robot teams), if all components share the same information, then it is possible to recreate essentially centralized decisions. Accordingly, true distributed architectures have costly communications and latent dissemination of information. This thrust emphasizes decision making in networked architectures through issues such as understanding efficient network topologies and determining what and when to communicate.

- *Evolving and uncertain environments:* An appeal of distributed learning is the potential to adapt to evolving conditions and react to disruptions such as failed components or adversarial actions. This thrust emphasizes issues of robustness and adaptability in distributed decision architectures as well as learning about the unknown environment.

Our objective throughout is to understand the theoretical limits of achievable performance imposed by distributed decision architectures and to derive constructive and computationally feasible implementation algorithms.

2.1 Distributed online resource allocation

ADMM-based methods for Multi-agent Optimization over Undirected and Directed Networks: We consider a multi-agent optimization problem: a network of agents is solving a global optimization problem over a network where the objective function is given by the sum of privately known convex local objective functions. For this problem, we present a fully distributed Alternating Direction Method of Multipliers (ADMM) based method. We show that this method converges at the rate $O(1/T)$ (T is the iteration number), which is much faster than the $O(1/\sqrt{T})$ convergence rate of commonly used subgradient based methods. We study the dependence of the algorithm on the network structure. We also develop asynchronous implementations of this algorithm with similar performance. In more recent work, we developed new ADMM type algorithms that use broadcast-based communication among agents (as opposed to pairwise communication between pairs of agents connected through an edge), and hence are more efficient in terms of communication exchanges. In recent work, we have developed new subgradient and ADMM-type algorithms that can be implemented in a distributed manner over directed networks. The key idea of these algorithms is to update estimates at each node based on incoming information, but scale the information of each node by the "right weights" such that in the limit the graph gets balanced, i.e., at each node the amount of incoming information is equal to the amount of outgoing information.

Incremental Methods for Additive Cost Convex Optimization. Motivated by machine learning problems over large data sets and distributed optimization over networks, we consider the problem of minimizing the sum of a large number of convex component functions. We study incremental methods for solving such problems that use information about a single component function at each iteration. In particular, we study first order incremental methods and provide new rate guarantees for such methods under smoothness assumptions. We provide lower bounds, which show that these methods cannot converge faster than rate $1/k$, where k is the number of iterations. We also develop second order incremental methods and provide convergence rate results in which condition number dependence is removed. In more recent work, we consider incremental aggregated gradient methods, which process a single component function, but keep a memory of the most recent gradients of all component functions to construct an approximation of the full gradient. In contrast with incremental gradient methods, we show that incremental aggregated gradient methods converge linearly.

Conditions for the existence (and methods for the design) of order-constrained distributed estimators: Consider that an autonomous LTI plant is given and that a network of local observers assesses the output vector of the plant. Each local observer in the network is LTI and measures a

portion of the output vector. Each local observer computes a state estimate using its measurements and the state estimates of other local observers shared according to a pre-selected neighborhood structure. We have proposed an update rule that is a natural generalization of consensus, and for which we determine necessary and sufficient conditions for the existence of parameters that lead to asymptotic omniscience of the state of the plant at all local observers. In comparison with our prior work, during this reporting period, we obtained necessary and sufficient conditions that are less stringent and we developed a design method that guarantees that the average dimension of the state of the observers does not exceed that of the plant plus one. We have also shown that our results are a natural extension of existing consensus methods to the dynamic case. A most general version of this work [NM3] is under review for publication in the IEEE Transactions on Automatic Control.

Design of event-based optimal remote estimation systems: We have proposed two new formulations to study the design of optimal remote estimation systems in which multiple agents have unequal measurements but wish to estimate a quantity that may depend on all measurements. In the first formulation [NM6] information is conveyed from the sensors to a fusion center via an action dependent channel, which can be used to account for an energy harvesting mechanism or a human assistant. The case in which each transmission incurs a cost is analyzed in [NM2]. In the second formulation, sensors must communicate with a fusion center via a collision channel, which models the effect of interference caused by simultaneous transmissions. Here we propose a new method based on constrained moment optimization to characterize and compute optimal solutions [NM5].

Distributed methods for multi-agent optimization over networks. Many of today's data processing takes place over large-scale networks in which information is gathered and processed locally. A prominent example is a machine learning application where a network of agents aims to estimate a model (or a parameter) to optimize a global loss function using decentralized computations based on locally collected data. This has motivated a recent burgeoning literature on designing decentralized optimization algorithms for efficient processing of this local information. These problems are typically formulated as a global optimization problem (referred to as a multi-agent optimization problem) where the objective function is the sum of local objective functions of agents that are connected through a network. Many of the existing algorithms in the literature are based on first-order subgradient methods, which have slow convergence rates (given by $O(1/\sqrt{k})$ where k is the iteration number), making them impractical in many large scale applications.

Last year, under the auspices of this project, we developed asynchronous decentralized algorithms based on the classical Alternating Direction Method of Multipliers (ADMM) for solving multi-agent optimization problems. In recent work, we developed new distributed ADMM algorithms that involve edge-based or broadcast based communication among the agents. We show that both the objective function values and the feasibility violation converge with the faster rate $O(1/k)$. We study the dependence of the convergence rate on the underlying network structure. In particular, we provide bounds on the objective function value improvement and feasibility violation as a function of the number of edges, maximum degree in the network and the singular values of the communication matrix.

Almost all the work in the literature on distributed optimization assumes that the underlying

network is symmetric implying that the communication among nodes is bidirectional. This assumption does not hold when nodes have different local interference patterns and different power transmission levels. In recent work, we developed a distributed ADMM-type algorithms for multi agent optimization over directed networks. These algorithms involve each agent updating primal and dual variables based on information received through incoming links.[?] We normalize the information entering a node with the right weight so that in the limit the network becomes balanced (i.e., amount of information entering and leaving a node is equalized), and the agent estimates converge to an optimum solution of the global optimization problem.

Distributed optimization with equality constraints. We developed and analyzed two discrete-time, distributed optimization algorithms executed by a set of agents whose interactions are subject to a communication graph. The algorithms can be applied to optimization problems where the cost function is expressed as a sum of functions, and where each function is associated to an agent. In addition, the agents can have equality constraints as well. The algorithms are not consensus-based and can be applied to non-convex optimization problems with equality constraints. We demonstrated that the first distributed algorithm results naturally from applying a first order method to solve the first order necessary conditions for a lifted optimization problem with equality constraints; the solution of our original problem is embedded in the solution of this lifted optimization problem. Using an augmented Lagrangian idea, we derived a second distributed algorithm that requires weaker conditions for local convergence compared to the first algorithm. For both algorithms we address the local convergence properties.

Distributed method of multipliers for nonlinear programming. We considered a distributed optimization problem, where a set of agents interacting through a communication graph have as common goal the minimization of a function expressed as a sum of (possibly non-convex) differentiable functions. Each function in the sum corresponds to an agent and each agent has associated an equality constraint. We investigated how the standard method of multipliers can be used to solve an optimization problem with equality constraints in a distributed manner. The method of multipliers was applied to a lifted optimization problem whose solution embeds the solution of the original problem. We modified the standard convergence results to deal with the fact that the (local) minimizers of the lifted optimization problem are not regular, as a result of the distributed formulation.

Fast convergence in general networks. Suppose that agents are located at the nodes of a network and they update their actions asynchronously by choosing noisy best responses to their neighbors' actions. We show that when the level of noise lies in an intermediate range, the convergence time scales linearly in the number of agents, and is bounded uniformly irrespective of the topological characteristics of the network, such as the degree of clustering, diameter, and path length. The methods of proof rely on a combination of stochastic dynamical systems theory and martingale theory. Analogous results hold when the agents experience heterogeneous payoff shocks. This work was recently published in the *Proceedings of the National Academy of Sciences*.

Fast convergence in bilateral networks. An important class of networks in applications are those in which one class of agents (e.g., firms) needs to be matched with another class of agents

(e.g., workers) so as to maximize total payoffs. There are standard algorithms for solving such assignment problems in polynomial time, but they require full information about the value of all potential matches, which in practice may not be available. We show that there exist simple distributed learning algorithms that achieve the same result even when agents have no information about the payoffs of their potential partners. The proof uses novel techniques characterizing random walks on trees. The result has potential application to the design of clearing algorithms for on-line matching markets, which are typically highly decentralized and involve huge numbers of agents.

Convergence rates in weakly acyclic games. Previous research has shown that there exist fully distributed learning algorithms that converge rapidly to equilibrium for congestion games, but to date little is known about convergence rates in other classes of weakly acyclic games, such as coordination games, supermodular games, and dominance solvable games. We show that weak acyclicity is not sufficient to guarantee fast convergence; in fact there exist very simple weakly acyclic games such that virtually any distributed payoff-based learning algorithm takes exponentially long (in the population size) to come close to Nash equilibrium. Paradoxically, however, if the system is subjected to random shocks that temporarily block communications, so that agents are sometimes unable to determine which strategies yield higher payoffs, the learning process can be greatly accelerated. In fact, if the game is generic and weakly acyclic then for a large class of shock distributions and distributed learning rules, the expected waiting time to come ϵ -close to Nash equilibrium is bounded by a polynomial in $1/\epsilon$ that is independent of the population size N .

Achieving system optimality with fully distributed learning. We demonstrate a simple payoff-based learning rule that is completely decentralized, and that leads to an efficient configuration of actions in any n -person finite strategic-form game with generic payoffs. The algorithm requires no communication: agents respond solely to changes in their own realized payoffs, which are affected by the actions of other agents in the system in ways that they do not necessarily understand. The method can be applied to the optimization of complex systems with many distributed components, such as the routing of information in networks and the design and control of wind farms. The proof of the proposed learning algorithm relies on the theory of large deviations for perturbed Markov chains.

Multi-agent welfare optimization over continuous sets. Related to the above, we address the more challenging, continuous action space analog of these results. We investigated a distributed, model-free optimization problem in the context of multi-agent systems. The set-up comprises of a fixed number of agents, each of which can pick an action and receive/measure a private utility function that can depend on the collective actions taken by all agents. The exact functional form (or model) of the agent utility functions is unknown, and an agent can only measure the numeric value of its utility. The objective of the multiagent system is to optimize the welfare function (i.e. sum of the individual utility functions). A model-free, distributed, on-line learning algorithm was developed that achieves this objective. Our proposed solution requires information exchange between the agents over an undirected, connected communication graph, and is based on ideas from extremum seeking control. We also proved a result on local convergence of the proposed algorithm to an arbitrarily small neighborhood of a local minimizer of the welfare function.

New Notions of Stability of Evolutionary Games. In a collaboration between Martins and Shamma, we are exploring new notions of solution of an evolutionary game for cases when certain passivity conditions hold. We have made some preliminary progress reported in [NM2], where we propose the new notions of integrability and non-positive correlation and show that they have computational advantages and are equivalent to passivity. This will allow us to study new types of population games, such as when there is a delay between the state and the corresponding payoff.

Active Learning and Best-Response Dynamics. We examine an important setting for engineered systems in which low-power distributed sensors are each making highly noisy measurements of some unknown target function (such as whether there is a high or low concentration of a dangerous chemical at their location). A center wants to accurately learn this function by querying a small number of sensors, which ordinarily would be impossible due to the high noise rate. The question we address is whether local communication among sensors, together with natural best-response dynamics in an appropriately-defined game, can denoise the system without destroying the true signal and allow the center to succeed from only a small number of active queries. By using techniques from game theory, machine learning, and empirical processes, we prove positive (and negative) results on the denoising power of several natural dynamics. We then show how we can use our recent advances in active learning/querying in the presence of noise (aka agnostic active learning), in order to provide provable guarantees for the center to learn a good approximation to the target function from a small number of queries to the agents. We also show experimentally that this process can achieve low error from very few queries, performing substantially better than active or passive learning without first using the game theoretic denoising dynamics, as well as passive learning with denoising.

Improved Distributed Principal Component Analysis. We study the distributed computing setting in which there are multiple servers, each holding a set of points, who wish to compute functions on the union of their point sets. A key task in this setting is Principal Component Analysis (PCA), in which the servers would like to compute a low dimensional subspace capturing as much of the variance of the union of their point sets as possible. Given a procedure for approximate PCA, one can use it to approximately solve a variety of problems such as k-means clustering and low rank approximation. The essential properties of an approximate distributed PCA algorithm are its communication cost and computational efficiency for a given desired accuracy in downstream applications. In this work, we provide new algorithms and analyses for distributed PCA which lead to improved communication and computational costs for k-means clustering and related problems. Our empirical study on real world data shows a speedup of orders of magnitude, preserving communication with only a negligible degradation in solution quality.

2.2 Networked operations

Coevolutionary Model of Strategic Network Formation: In foundational models of network formation, the mechanisms for link formation are based solely on network topology. For example, preferential attachment uses degree distributions, whereas a strategic connections model uses inter-node distances. These dynamics implicitly presume that such benefits and costs are instantaneous functions of the network topology. A more detailed model would include that benefits and

costs are themselves derived through a dynamic process, which, in the absence of time-scale separation, necessitates a co-evolutionary analysis. This work introduces a new co-evolutionary model of strategic network formation. In this model, network formation evolves along with the flow of benefits from one node to another. We examine the emergent equilibria of this combined dynamics of network formation and benefit flow. We show that the class of strict equilibria is stable (or robust to small perturbations in the benefits flows).

Game theoretic self-reconfiguration: This work formulates the homogeneous two- and three-dimensional self-reconfiguration problem over discrete grids as a constrained potential game. We develop a game-theoretic learning algorithm based on the Metropolis-Hastings algorithm that solves the self-reconfiguration problem in a globally optimal fashion. Both a centralized and a fully distributed algorithm are presented and we show that the only stochastically stable state is the potential function maximizer, i.e. the desired target configuration. These algorithms compute transition probabilities in such a way that even though each agent acts in a self-interested way, the overall collective goal of self-reconfiguration is achieved. Simulation results confirm the feasibility of our approach and show convergence to desired target configurations.

Consensus Algorithms and Applications to Distributed Optimization Distributed consensus in networks of interacting agents has been revisited with a view toward stronger results than currently available, as well as new insights into various issues such as consensus prediction and intrusion detection. Among the achievements in this research topic is a new algorithm for distributed optimization on a network that also includes development of a significantly faster consensus algorithm than the classical algorithm based on simple averaging of neighbors' opinions at each time step. Thus, consider a distributed optimization problem in which a network of agents wants to minimize the sum of their individual cost functions. Two fast algorithms based on a finite-time consensus protocol are proposed. The first algorithm is of subgradient-type and assumes that the individual cost functions are either convex with bounded subgradients, or differentiable such that their sum is convex and has a Lipschitz continuous gradient. In the former case, convergence to a solution neighborhood is proved, while in the latter asymptotic convergence can be ensured when using a constant step-size and even when some of the individual cost functions are non-convex. Moreover, the convergence rate of the new algorithm is similar to that of centralized subgradient methods. With slight modifications, the algorithm is shown to work even when the network weight matrix is only row-stochastic. When the optimization problem is unconstrained and the local cost functions are quadratic, by combining the ratio-consensus protocol and the finite-time consensus protocol, another algorithm is proposed that converges in finite time. Simulation examples are given to illustrate the algorithms, showing rapid convergence to the optimal solution after which the solution is maintained for future time. This algorithm is also used to determine an observation sub-network of a given network of agents, where monitoring the agents of the sub-network is used to predict the ultimate consensus behavior of the overall network.

Modeling of Networked Communities Community-based Susceptible-Infected-Recovered (SIR) and Susceptible-Infected-Susceptible (SIS) models of infection/innovation diffusion are introduced for heterogeneous social networks in which agents are viewed as belonging to one of a finite number of communities. Agents are assumed to have well-mixed interactions within and between

communities. The communities are connected through a backbone graph which defines an overall network structure in the models. The models are used to determine conditions for outbreak of an incipient infection. The role of the strengths of the connections between communities in development of an outbreak is studied. Long term behavior of the diffusion is also studied. Percolation theory is also brought to bear on these questions as an independent approach separate from this dynamic multi-community modeling approach, and results obtained using both approaches are compared and found to be in agreement in the limit of infinitely large populations within each community. Based on the proposed models, three classes of marketing problems are formulated and studied: referral marketing, seeding marketing and dynamic marketing. It is found that referral marketing is simple because it can be formulated in terms of convex optimization. Also, both seeding marketing and dynamic marketing are shown to enjoy a useful property, namely "continuous monotone submodularity." Based on this property, a greedy heuristic is proposed which yields solutions with approximation ratio no less than $1-1/e$. Also, dynamic marketing for SIS models is reformulated into an equivalent convex optimization to obtain an optimal solution. Both cost minimization and trade-off of cost and profit are analyzed. Next, the proposed modeling framework is applied to the problem of competition of multiple companies in marketing of similar products. Two classes of such marketing problems are considered, namely marketing of durable consumer goods (DCG) and fast-moving consumer goods (FMCG). It is shown that an epsilon-equilibrium exists in the DCG marketing game and a pure Nash equilibrium exists in the FMCG marketing game. The Price of Anarchy (PoA) in both marketing games is found to be bounded by 2. Also, it is shown that any two Nash equilibria for the FMCG marketing game agree almost everywhere, and a distributed algorithm converging to the Nash equilibrium is designed for the FMCG marketing game. Finally, a preliminary investigation was carried out to explore possible concepts of dynamic network centrality for diffusions. Standard network centrality concepts tend to rank the importance of nodes from a static point of view. For example, Page Rank is based on a computation of the power of a node in terms of its level of connection to the rest of the network through page links. In a dynamic social network, the power of a node should reflect the influence that the node has over the network over time. Among the preliminary observations in this work, it is found that if an infection does not break out, diffusion centrality is closely related to Katz centrality; when an infection does break out, diffusion centrality is found to be closely related to eigenvector centrality.

Distributed state omniscience. Consider that an autonomous LTI plant is given and that a network of local observers assesses the output vector of the plant. Each local observer in the network is LTI and measures a portion of the output vector. Each local observer computes a state estimate using its measurements and the state estimates of other local observers shared according to a pre-selected neighborhood structure. We have proposed an update rule that is a natural generalization of consensus, and for which we determine necessary and sufficient conditions for the existence of parameters that lead to asymptotic omniscience of the state of the plant at all local observers. In comparison with our prior work, over the last year we obtained necessary and sufficient conditions that are less stringent. In particular, we showed that if the conditions hold then there is a systematic design method to obtain a solution, and no solution exists otherwise. In fact, if the condition fails then no scheme can attain asymptotic omniscience even when nonlinear and time-varying configurations are allowed. We have extended our previous results to the case when the network is directed. We have also found necessary and sufficient conditions for the existence of an

omniscience scheme, based on properties of the network. We have also shown that our results are a natural extension of existing consensus methods to the dynamic case. This grant also partially supported a MS student who developed planning algorithms that we expect to use in conjunction with our distributed omniscience algorithms to achieve the coordination of multiagent systems. The student has concluded his MS Thesis and is now employed at GE research center at NY.

Person-by-person optimality of remote estimation policies with communication costs. Recently we were able to characterize for the first time the person-by-person optimal solutions to the problem proposed in Lipsa and Martins, “Remote State Estimation with Communication Costs for First-Order LTI Systems,”, in which an optimal estimation problem is posed where costly transmissions are included as part of the optimization criterion. Extending the prior work, the source is now an LTI system of arbitrary order. We also formulate algorithms that are numerically efficient.

Collaborative coordination of UAVs in the NAS with safety guarantees. We developed new fundamental methodologies for high performance and provably safe autonomous and collaborative control and operation of autonomous unmanned aerial systems (UAS) in the national airspace (NAS), where both UAS and manned aircraft fly. The proposed framework [3] is model-based, and emphasizes multiple scales in time and space as well as hybrid systems mathematics to capture both the analog and logical components of control functionalities. We developed on-line control laws that allow for multiple UAS agents to reconfigure to different formations with proofs of safety and convergence while navigating in integrated airspace with piloted air vehicles. We demonstrated our results in simulated scenaria with both cooperative and uncooperative air vehicles.

Distributed opportunistic scheduling for ad-hoc networks. In heavily constrained communication environments in multi-agent systems, agents must learn to identify opportunities for communication and patterns of such opportunities. This leads to opportunistic (real-time) scheduling problems for communication between the agents. The topic of opportunistic scheduling for wireless ad-hoc networks has been studied for single-channel networks in several recent works. Since now many wireless systems provide multiple channels for data transmission, the opportunistic scheduling problem is of practical interest for networks with multiple channels. We investigated a distributed opportunistic scheduling problem to exploit the channel fluctuations in wireless ad-hoc networks. In this problem, channel probing is followed by a transmission scheduling procedure that is executed independently within each link in the network. We investigated this problem for the popular block-fading channel model, where channel dependencies are inevitable between different time instances during the channel probing phase. Different from existing works, we explicitly considered this type of channel dependencies and its impact on the transmission scheduling and hence the system performance. We used optimal stopping theory to formulate this problem, but at carefully chosen time instances at which effective decisions are made. The problem can then be solved by a new stopping rule problem where the observations are independent between different time instances. Since the stopping rule problem has an implicit horizon determined by the network size, we first characterized the system performance using backward induction. We developed one recursive approach to solve the problem and showed that the computational complexity is linear with respect to the network size. Due to its computational complexity, we presented an approximation for performance analysis and developed a metric to check how good the approximation

is. We characterized the achievable system performance if we ignore the finite horizon constraint and applied the stopping rules based on the infinite horizon analysis nevertheless. We presented an improved protocol to reduce the probing costs which requires no additional cost. We characterized the performance improvement and the energy savings in terms of the probing signals. We obtained numerical results based on our mathematical analysis with various settings of parameters.

Distributed temporal coordination. An important aspect of distributed decision problems is temporal coordination. We investigate this issue in the setting of multi-agent sequential hypothesis. Two agents measure private information and must make a one-time declaration about the state of world after costly gathering information. There is a penalty for both incorrect conclusions and a temporal mismatch in making the declaration. The associated Bayes risk functions explicitly incorporate costs of taking private/public measurements, costs of time-difference, disagreement in actions of agents, and costs of false declaration/choices about the state of nature. We show that the corresponding sequential decision processes have well-defined value functions with respect to (a) the belief states for the case of conditional independent private noisy measurements that are also assumed to be independent identically distributed over time, and (b) the information states for the case of correlated private noisy measurements.

2.3 Evolving and uncertain environments

Scalable asynchronous multiagent decision-making under uncertainty Decentralized partially observable Markov decision processes (Dec-POMDPs) can model rich environments for cooperative decision-making under uncertainty, but they typically require synchronous decision-making: every agent determines an action to execute, and then executes it in a single time step. This restriction is problematic for two reasons. First, many systems have a set of controllers (for e.g., waypoint navigation, grasping an object, waiting for a signal), and planning consists of sequencing the execution of those controllers. These controllers are likely to require different amounts of time, resulting in suboptimality or impossible execution when using synchronous decision-making. Second, the planning complexity of a Dec-POMDP is doubly exponential in the horizon. A planner that reasons about all of the agents' possible policies at every time step will only ever be able to make very short plans. To address asynchronous decision-making in these systems, we have extended the Dec-POMDP model to incorporate *macro-actions* (i.e., temporally extended actions which may require different amounts of time). Macro-actions enable the decision-making to take place at a higher level — at the level of deciding which macro-actions to execute — and the macro-actions themselves can then be executed to completion. By extending Dec-POMDP algorithms to the macro-action case, we showed that value that is similar to state-of-the-art Dec-POMDP approaches could be achieved on benchmark problems while using significantly less time. We also showed that small multi-robot warehousing problems could be solved that were orders of magnitude larger than problems solved by previous methods.

Exploiting structure in centralized planning for cooperative decentralized systems When decentralized systems are cooperative, offline planning can be completed in a centralized manner as long as the final policies that are generated can be executed in a decentralized manner. Last year, we showed that by using a centralized planning stage, but decentralized execution, we can

transform a Dec-POMDP into a *continuous-state MDP*. Using this novel transformation, we can then apply a wide range of powerful techniques for solving POMDPs and continuous-state MDPs. However, scalability remains limited when the number of agents or problem variables becomes large. Now, we have shown that, under certain separability conditions of the optimal value function, the scalability of this approach can increase considerably. This separability is present when there is locality of interaction, which — as other approaches have already shown — can be exploited to improve performance. Locality of interaction is common in multiagent systems as agents are likely to interact not with all other agents in the system, but only some local subset. Unlike most previous methods, our novel continuous-state MDP algorithm retains optimality and convergence guarantees. Results show that the extension using separability can scale to a large number of agents and domain variables while maintaining optimality.

Network Flow under Cascaded Failures We propose a dynamical model for cascading failures in single-commodity network flows. In the proposed model, the network state consists of flows and activation status of the links. Network dynamics is determined by a, possibly state-dependent and adversarial, disturbance process that reduces flow capacity on the links, and routing policies at the nodes that have access to the network state, but are oblivious to the presence of disturbance. Under the proposed dynamics, a link becomes irreversibly inactive either due to overload condition on itself or on all of its immediate downstream links. The coupling between link activation and flow dynamics implies that links to become inactive successively are not necessarily adjacent to each other, and hence the pattern of cascading failure under our model is qualitatively different than standard cascade models. The magnitude of a disturbance process is defined as the sum of cumulative capacity reductions across time and links of the network, and the margin of resilience of the network is defined as the infimum over the magnitude of all disturbance processes under which the links at the origin node become inactive. We propose an algorithm to compute an upper bound on the margin of resilience for the setting where the routing policy only has access to information about the local state of the network. For the limiting case when the routing policies update their action as fast as network dynamics, we identify sufficient conditions on network parameters under which the upper bound is tight under an appropriate routing policy. Our analysis relies on making connections between network parameters and monotonicity in network state evolution under proposed dynamics.

Unified theory of Cascades We study a model for cascade effects over finite networks based on a deterministic binary linear threshold model. Our starting point is a networked coordination game where each agents payoff is the sum of the payoffs coming from pairwise interaction with each of the neighbors. We first establish that the best response dynamics in this networked game is equivalent to the linear threshold dynamics with heterogeneous thresholds over the agents. While the previous literature has studied such linear threshold models under the assumption that each agent may change actions at most once, a study of best response dynamics in such networked games necessitates an analysis that allows for multiple switches in actions. In this paper, we develop such an analysis. We first establish that agent behavior cycles among different actions in the limit, we characterize the length of such limit cycles, and reveal bounds on the time steps required to reach them. We finally propose a measure of network resilience that captures the nature of the involved dynamics. We prove bounds and investigate the resilience of different network

structures under this measure.

Network tomography. In uncertain and evolving networked systems environments, an important problem is that of estimating and inferring topological, link and node properties of the dynamic graph, based on adaptive measurements from only a subset of the nodes (which can be dynamically varying itself). This is the celebrated Network Tomography problem, which has attracted tremendous attention in various forms, both discrete and continuous. During the reporting period, and building on our previous mathematical work on these problems, we investigated [3] several rigorously defined network tomography problems from applications ranging from communication networks, to social networks. The universal abstraction we developed [3] involves the inference of various network structural and parametric properties from observations of certain “probing” processes from a subset of network nodes which we call the “boundary nodes” of the network. We showed that these problems lead to mathematical problems of “deconvolution” over unconventional semirings, inversion of integrals over trees of the underlying graph, Radon transform over symmetric spaces, and completion of sparse matrices albeit in non-conventional semirings. Further we illustrated our algorithms and results in several applied problems. Remarkably, some of our problem formulations and solutions are inspired by generalizations of “electrical impedance problems” in various non-nonconventional spaces. Some of our solutions and methods have a variational interpretation.

Influence Function Learning in Information Diffusion Networks. Can we learn the influence of a set of people in a social network from cascades of information diffusion? This question is often addressed by a two-stage approach: first learn a diffusion model, and then calculate the influence based on the learned model. Thus, the success of this approach relies heavily on the correctness of the diffusion model which is hard to verify for real world data. In this work, we exploit the insight that the influence functions in many diffusion models are coverage functions (a special interesting case of submodular functions), and propose a novel parameterization of such functions using a convex combination of random basis functions. Moreover, we propose an efficient maximum likelihood based algorithm to learn such functions directly from cascade data, and hence bypass the need to specify a particular diffusion model in advance. We provide both theoretical and empirical analysis for our approach, showing that the proposed approach can provably learn the influence function with low sample complexity, be robust to the unknown diffusion models, and significantly outperform existing approaches in both synthetic and real world data.

Learning Economic Parameters from Revealed Preferences. We study the classic “revealed preference” setting, where we assume that agents, facing prices, will choose to buy the bundle of goods that they most prefer among all bundles that they can afford, according to some concave, non-decreasing utility function. The goal is to produce a model of the agents’ utility function that can explain her behavior based on past data. Work on this topic has a long history in economics, beginning with the seminal work by Samuelson (1948). Traditionally, this work has focused on the “rationalization” or “fitting the sample” problem, in which explanatory utility functions are constructively generated from finitely many agent price/purchase observations. Note, however, that just because a function agrees with a set of data does *not* imply that it will necessarily predict future purchases well.

A recent exciting line of work, starting with Beigman and Vohra introduced a statistical learning analysis of the problem of learning the utility function from past data with the explicit formal goal of having predictive or forecasting properties (provably). In this work, we advance this line of work by providing sample complexity guarantees and efficient algorithms for a number of important classes. By drawing a connection to recent advances in multi-class learning and by casting these problems as “structured prediction” problems we obtain the right sample complexity and computationally efficient algorithms for a number of interesting cases (essentially solving a number of open questions in the earlier work).

Fast Convergence in Semi-Anonymous Potential Games. The log-linear learning algorithm has been extensively studied in both the game theoretic and distributed control literature. One of the central appeals with regards to log-linear learning for distributed control is that it often guarantees that the agents behavior will converge in probability to the optimal configuration. However, one of the central issues with log-linear learning for this purpose is that the worst case convergence time can be prohibitively long, e.g., exponential in the number of players. In this work, we formalize a modified log-linear learning algorithm that exhibits a worst case convergence time that is roughly linear in the number of players. We prove this characterization in semi-anonymous potential games with limited populations. That is, potential games where the agents utility functions can be expressed as a function of aggregate behavior within each population.

Influencing Social Behavior in Networked Systems. Developing the infrastructure necessary to serve the needs of our community is a dominant engineering principle. A crucial design challenge associated with this task is ensuring that this infrastructure is utilized in an efficient manner. This challenge arises from the well-known fact that uninfluenced social systems can exhibit highly inefficient system behavior, e.g., transportation networks. Accordingly, the influence and coordination of social behavior is a fundamental challenge that engineers must account for. Here, we focus on the derivation of local mechanisms for coordinating social behavior in situations where the society’s sensitivity to such mechanisms is unknown. Preliminary work focuses on the use of tolling in congestion networks where the system design is uncertain of the population’s sensitivity to tolls.

Potential Games are Necessary to Ensure Pure Nash Equilibria in Cost Sharing Games. We consider the problem of designing distribution rules to share ‘welfare’ (cost or revenue) among individually strategic agents. There are many known distribution rules that guarantee the existence of a (pure) Nash equilibrium in this setting, e.g., the Shapley value and its weighted variants; however, a characterization of the space of distribution rules that guarantee the existence of a Nash equilibrium is unknown. Our work provides an exact characterization of this space for a specific class of scalable and separable games, which includes a variety of applications such as facility location, routing, network formation, and coverage games. Given arbitrary local welfare functions W , we prove that a distribution rule guarantees equilibrium existence for all games (i.e., all possible sets of resources, agent action sets, etc.) if and only if it is equivalent to a generalized weighted Shapley value on some ground welfare functions $W^?$, which can be distinct from W . However, if budget-balance is required in addition to the existence of a Nash equilibrium, then $W^?$ must be the same as W . We also provide an alternate characterization of this space in terms of generalized marginal

contributions, which is more appealing from the point of view of computational tractability. A possibly surprising consequence of our result is that, in order to guarantee equilibrium existence in all games with any fixed local welfare functions, it is necessary to work within the class of potential games.

Speed of convergence of distributed learning in Nash population games. Consider a finite n -person game G in strategic form. The Nash population game derived from G has N players in each of the n player positions of G , and the payoff to a given player is his expected payoff from playing a randomly drawn subset of players from the other positions. Thus the distribution of actions taken by the various players determines the payoff of each individual just as if these distributions were mixed strategies in the game G . We investigate the efficiency of distributed learning algorithms in this setting. Specifically, if the underlying game G is weakly acyclic, then so is the Nash population game derived from G . Given any such game and any initial state there exists a better reply path to a pure, strict Nash equilibrium. Therefore one might hope that at least some distributed better reply algorithms are efficient in the sense that the expected waiting time is polynomial in N . We show that this is not the case: there are quite simple games in which any distributed better reply algorithm takes exponentially long to come close to Nash equilibrium. However, if the learning process is subjected to aggregate shocks—say because information about others payoffs is temporarily interrupted—then the learning process is greatly accelerated. Indeed under suitable regularity conditions on the shock distributions and learning rules, for generic weakly acyclic games G the learning process comes ϵ -close to Nash equilibrium in an expected time that is bounded above for all N and is polynomial in $1/\epsilon$.

Distributed learning in online matching markets. Many modern markets are highly decentralized, involve huge numbers of agents, and are conducted in an on-line format. Examples include online markets for matching buyers and sellers of goods, workers and firms, hotels and clients, and so forth. In these environments the participants have very little information about the market as a whole or the preferences of the other agents. Thus the appropriate models of their behavior are completely uncoupled learning rules that depend only on experienced payoffs. We show that there is a very natural family of completely uncoupled adjustment mechanisms that lead to stable and optimal core allocations in matching markets of arbitrary size. Moreover the adjustment process puts highest probability on solutions that are equitable in the sense that matched agents split the surplus equally subject to the core constraints. A key point is that equitable outcomes are selected because they are more stable, not because the agents care about equity per se. The next stage of the research will be to characterize the efficiency of these learning algorithms.

Stochastic stability analysis in games with many agents. We develop a general model of stochastic learning in games and characterize the learning behavior in the large population limit over different time scales. Existing approaches obtain deterministic approximations of finite-horizon behavior in terms of the mean dynamics, and they describe infinite-horizon behavior and stochastic stability via limiting stationary distributions. To complement these analyses, and to explain the relations between them, we provide a large deviations analysis that describes the probabilities of excursions against the flow of the mean dynamic over intermediate time scales. This “local” large deviations analysis is combined with graph theoretic methods to obtain conclusions

about the “global” behavior of the process, in particular, the asymptotics of the stationary distribution and the expected times until exit from general domains. Detailed analyses are provided for two settings: potential games under the logit choice rule, and two-strategy games under arbitrary choice rules. In both cases, the large deviations analysis reveals that the processes are “reversible in the large”, providing a powerful new approach to understanding global behavior in these settings.

New Tools for Stochastic Stability Analysis. Emergent behavior in natural and manmade systems can often be characterized by the limiting distribution of a special class of Markov processes termed regular perturbed processes. Resistance trees have gained popularity as a computationally efficient way to characterize the stochastically stable states (i.e., support of the limiting distribution); however, there are three main limitations of that approach. First, it often requires finding a minimum weight spanning tree for each state in a potentially large state space. Second, perturbations to transition probabilities must decay at an exponentially smooth rate. Lastly, that approach is shown to hold purely in the context of finite Markov chains. Here, we focus on addressing these limitations by developing new tools for characterizing the stochastically stable states. First, we provide necessary conditions for stochastic stability via a coarse, and less computationally intensive, state space analysis. Next, we identify necessary conditions for stochastic stability when smooth convergence requirements are relaxed. Lastly, we establish similar tools for stochastic stability analysis in Markov chains over a continuous state space.

Robustness of stochastic stability. The notion of stochastic stability is used in game theoretic learning to characterize which joint actions of players exhibit high probabilities of occurrence in the long run. This paper examines the impact of two types of errors on stochastic stability: i) small unstructured uncertainty in the game parameters and ii) slow time variations of the game parameters. In the first case, we derive a continuity result bounds the effects of small uncertainties. In the second case, we show that game play tracks drifting stochastically stable states under sufficiently slow time variations. The analysis is in terms of Markov chains and hence is applicable to a variety of game theoretic learning rules. Nonetheless, the approach is illustrated on the widely studied rule of log-linear learning. The results are applied in both simulation and laboratory experiments to distributed area coverage with mobile robots.

Convergence Guarantees for a Decentralized Algorithm Achieving Pareto Optimality. Distributed learning is an important component in distributed optimization decision making systems using local information. Simple such systems achieving excellent performance with as few as possible assumptions are of particular interest. We considered N agents, each picking actions from a finite set and receiving a payoff according to its individual utility function that may depend on the actions picked by others. An agent has no knowledge about the functional form of its utility and can only measure its instantaneous value. We assumed that all agents pick actions and receive payoffs synchronously. For this setting, a fully decentralized iterative algorithm for achieving Pareto optimality i.e. picking actions that maximize the sum of all utilities was recently proposed by Marden et. al., from our team, that lacks convergence guarantees. By scheduling a certain noise parameter to go to zero along iterations of this algorithm, we derived conditions that guarantee convergence in probability. The algorithm utilizes indirect communications between the agents resulting from the impact of an agents actions on the others. We improved this algorithm in our most recent work to

incorporate exchange of certain bit-valued information between the agents over a directed communication graph. We introduced the notion of an interaction graph is to encode known interactions in the system. We eliminated the restrictions on the payoff structure and derived conditions that guarantee convergence to welfare minimizing actions w.p. 1, under the assumption that the union of the interaction graph and the communication graph is strongly connected.

Learning behaviors for coordination in networked stochastic systems. We analyzed the effects of inter-agent behavior learning for coordination in networked systems in the framework previously developed and reported. The agents have to make a decision on whether to cooperate or not in a group effort based on their understanding of their neighbors behaviors. When making a decision, an agent is influenced by its knowledge about others behaviors. Agents understanding of others behaviors is shaped through observing their actions over a long time. We have modeled the decision making on whether to cooperate in a group effort as a result of a series of two-person games between agents, where the payoff of each agent is computed as the sum of its payoffs from each of these games. The agents initially have different behaviors. In order to maximize their payoff, they need to learn the others behavior and coordinate with them. We considered a behavior learning algorithm for a class of behavior functions and studied its effects on the emergence of coordination in the network. The conditions under which the learning algorithm converges were studied. We showed that for a class of piecewise linear functions the learning algorithm results in an extension of non-homogeneous consensus protocol to the more general case of block-stochastic matrices.

Selfish Response to Epidemic Propagation. Spreading information via various means in a system of networked agents has substantial effects on their decision making. Learning is an important aspect of such problems. An epidemic that spreads in a network calls for a decision on the part of the network users. They have to decide whether to protect themselves or not. Their decision depends on the trade-off between the perceived infection and the protection cost. Aiming to help users reach an informed decision, security advisories provide periodic information about the infection level in the network. We studied the best-response dynamic in a network whose users repeatedly activate or de-activate security, depending on what they learn about the infection level. Our main result is the counterintuitive fact that the equilibrium level of infection increases as the users learning rate increases. The same is true when the users follow smooth best-response dynamics, or any other continuous response function that implies higher probability of protection when learning a higher level of infection. In both cases, we characterized the stability and the domains of attraction of the equilibrium points. Our finding also holds when the epidemic propagation is simulated on human contact traces, both when all users are of the same best-response behavior type and when they are of two distinct behavior types.

Exploiting Ontology Structures and Unlabeled Data for Learning. In a number of modern applications of machine learning (including text understanding, medical diagnosis, search, and vision), one would like to learn multiple related tasks from primarily, or even solely, unlabeled data. One approach to this task is to take advantage of known relationships among the classes being learned, which are often captured by an ontology; for instance, this ontology might say that an example to be classified cannot be both a company and a product, or that if someone is an athlete,

then they are also a person. Such an ontology can either be provided by domain experts or learned from past data. In this work, we develop and analyze a theoretical model aimed at providing a mathematical understanding of learning methods of this type. We present both information-theoretic results as well as efficient algorithms. We show in this model that an ontology, which species the relationships between multiple outputs, in some cases is sufficient to completely learn a classification using a large unlabeled data source.

Active and passive learning of linear separators under log-concave distributions. We analyze active learning algorithms, which only receive the classifications of examples when they ask for them, and traditional passive (PAC) learning algorithms, which receive classifications for all training examples. We provide new results concerning label efficient, polynomial time, passive and active learning of linear separators under logconcave and more generally nearly log-concave distributions. We prove that active learning provides an exponential improvement over PAC (passive) learning of homogeneous linear separators under nearly log-concave distributions. Building on this, we provide a computationally efficient PAC algorithm with optimal (up to a constant factor) sample complexity for such problems. This resolves an open question of concerning the sample complexity of efficient PAC algorithms under the uniform distribution in the unit ball. Moreover, it provides the first bound for a polynomial-time PAC algorithm that is tight for an interesting infinite class of hypothesis functions under a general and natural class of data-distributions, providing significant progress towards a longstanding open question of (Ehrenfeucht et al.,1989; Blumer et al.,1989). We also provide new bounds for active and passive learning in the case that the data might not be linearly separable, both in the agnostic case and under the Tsybakov low-noise condition. To derive our results, we provide new structural results for (nearly) log-concave distributions, which might be of independent interest as well.

Programmable self-assembly and self-reconfiguration. In distributed self-assembly, a multitude of nodes, or agents, seek to form copies of a given target assembly. Nodes encounter each other in spontaneous pairwise interactions and decide whether or not to form or sever edges based on a fixed set of local interaction rules described by a graph grammar. In contrast to existing results, our work insist on especially stringent locality requirements. Agents can only update their labels when they themselves form or sever an edge. Our main result is the introduction of a simple algorithm that achieves an asymptotically maximum yield in a probabilistic sense. We also consider the related problem of self-reconfiguration, in which a collection of modules have the ability to change their morphology, structure, and functionality through changing the relative position of the modules. We derive a novel approach for heterogeneous self-reconfiguration of a modular robot comprised of heterogeneous cubic modules. We show that the proposed self-reconfiguration algorithm can guarantee completion of heterogeneous self-reconfigurations by avoiding so-called hole obstructions.

On the $O(1/k)$ Convergence of Asynchronous Distributed Alternating Direction Method of Multipliers. Though there have been many important advances in the design of decentralized optimization algorithms for multi-agent optimization problems, several challenges still remain. First, many of these algorithms are based on first-order subgradient methods, which have slow convergence rates (given by $O(1/\sqrt{k})$ where k is the iteration number), making them impractical

in many large scale applications. Second, with the exception of a few recent contributions, existing algorithms are synchronous, meaning that computations are simultaneously performed according to some global clock. In this work, we propose a fully asynchronous decentralized algorithm based on the classical Alternating Direction Method of Multipliers (ADMM) for solving a global optimization problem, where the objective function is the sum of privately known local objective functions of the agents and the decision variables are coupled via linear constraints. We adopt the following asynchronous implementation for our algorithm: at each iteration k , a random subset of the constraints are selected, which in turn selects the components of the decision vector that appear in these constraints. We refer to the selected constraints as active constraints and selected components as active components (or agents). We design an ADMM-type primal-dual algorithm which at each iteration updates the primal variables using partial information about the problem data, in particular using cost functions corresponding to active components and active constraints, and updates the dual variables corresponding to active constraints. Our first result shows that the (primal) asynchronous iterates generated by this algorithm converge almost surely to an optimal solution. Our proof relies on relating the asynchronous iterates to full-information iterates that would be generated by the algorithm that use full information about the cost functions and constraints at each iteration. In particular, we introduce a weighted norm where the weights are given by the inverse of the probabilities with which the constraints are activated and constructs a Lyapunov function for the asynchronous iterates using this weighted norm. Our second result establishes a performance guarantee of $O(1/k)$ for this algorithm, which to our knowledge is faster than the guarantees available in the literature for this problem. More specifically, we show that the expected value of the difference of the objective function value and the optimal value as well as the expected feasibility violation converges to 0 at rate $O(1/k)$.

2.4 Networked operations

Youla-Like Parameterizations Subject to Quadratic-Invariance Constraints. Consider that a linear time-invariant (LTI) plant is given and that we wish to design a stabilizing controller for it. Admissible controllers are LTI and must belong to a pre-selected subspace that may impose structural restrictions, such as sparsity constraints. The subspace is assumed to be quadratically invariant (QI) with respect to the plant, which, from prior results, guarantees that there is a convex parametrization of all admissible stabilizing controllers provided that an initial admissible stable stabilizing controller is provided. This work addresses the previously unsolved problem of extending Youlas classical parametrization so that it admits QI subspace constraints on the controller. In contrast with prior parameterizations, the one proposed here does not require initialization and it does not require the existence of a stable stabilizing controller. The main idea is to cast the stabilizability constraint as an exact model-matching problem with stability restrictions, which can be tackled using existing methods. Furthermore, we show that, when it exists, the solution of the exact model-matching problem can be used to compute an admissible stabilizing controller. Applications of the proposed parametrization on the design of norm-optimal controllers via convex methods are also explored. Recently, we have extended our work developed in previous years, which was valid only for sparsity pattern constraints on the controller and now it admits more general subspace constraints.

Necessary and Sufficient Conditions for the Stabilizability of a Class of LTI Distributed Observers. Consider that an autonomous LTI plant is given and that a network of local observers assesses the output vector of the plant. Each local observer in the network is LTI and measures a portion of the output vector. Each local observer computes a state estimate using its measurements and the state estimates of other local observers shared according to a pre-selected neighborhood structure. We have proposed an update rule that is a natural generalization of consensus, and for which we determine necessary and sufficient conditions for the existence of parameters that lead to asymptotic omniscience of the state of the plant at all local observers. In comparison with our prior work, over the last year we obtained necessary and sufficient conditions that are less stringent. In particular, we showed that if the conditions hold then there is a systematic design method to obtain a solution, and no solution exists otherwise. In fact, if the condition fails then no scheme can attain asymptotic omniscience even when nonlinear and time-varying configurations are allowed.

Methods to design of distributed estimators in the presence of transmission constraints. Consider a distributed estimation network that consists of several transmitters/sensors and a receiver/fusion-center. Each sensor measures a random variable and the goal of the receiver is to use the information transmitted to it to compute a MMSE estimate of a given function of the variables. We have proposed a new framework where the transmission by the sensors is affected by collisions. Our preliminary results indicate that optimal strategies are event-driven of the threshold type, and that they may be asymmetrical even when all the parameters of the problem are symmetric. It can be shown that this new problem class is related not only to past results under this MURI but also to problems in optimal quantization and moment-matching.

Expander Families as Information Patterns for Distributed Control. Part of our research is focused on identifying best topologies for communication in multi-agent multi-tasking systems, including learning such topologies. We revisited the classical, and recently intensively studied, problem of the distributed control of a 1-D vehicle platoon. A state space formulation was developed with the main objective being to understand the effects of the underlying information exchange pattern, between the vehicles, on the control performance of the platoon. The symmetric control case was considered where each vehicle gives equal weight to all the information available to it in determining its control law. We showed that expander families of graphs when used as information patterns result in stability margins decaying to zero at rate of at most $O(1/N)$; an improvement over the previously known $O(1/N^2)$ decay with nearest neighbor type information patterns.

Social networks over wireless networks. Future communications consist of an increasing number of wireless parts, while simultaneously need to support the widespread multimedia applications imposed by social networks. These human-machine systems, driven by both real time social interactions (read agent task and intention interactions) and the challenges of the wireless networks design (read between agents communication network), call for efficient and easy to implement, distributed cross-layer algorithms for their operation. Performance metrics such as throughput, delay, trust, energy consumption, need to be improved and optimized aiming at high quality communications. As broadband wireless devices and networks are becoming ubiquitous these human-machine systems, that combine the social aspects and behavioral activities of humans with the technological

characteristics of the underlying physical networks, provide several important challenges in efforts to model them, evaluate their performance and dynamically control them so certain performance requirements are met. We developed novel models for these complex human-machine systems that incorporate social network behavioral models and wireless network models that are inspired from statistical physics (hyperbolic graphs). We investigated the performance of wireless network protocols that support and respond to the constraints implied by the social network they support. We investigated the coveted throughput-delay trade-off in static wireless multihop networks based on a computer-aided design of the backpressure scheduling/routing algorithm for networks embedded in hyperbolic space. Both routing and scheduling exploit the hyperbolic distances to orient the packets to the destination and prioritize the transmissions correspondingly. The proposed design provides us with the freedom of controlling its theoretical throughput optimality and of counterbalancing its practical performance through simulations, leading to significant improvements of the throughput-delay trade-off.

Distributed Opportunistic Scheduling for Wireless Ad-Hoc Networks with Block-Fading Model.

In heavily constrained communication environments in multi-agent systems, agents must learn to identify opportunities for communication and patterns of such opportunities. This leads to opportunistic (real-time) scheduling problems for communication between the agents. The topic of opportunistic scheduling for wireless ad-hoc networks has been studied for single-channel networks in several recent works. Since now many wireless systems provide multiple channels for data transmission, the opportunistic scheduling problem is of practical interest for networks with multiple channels. We studied the problem of opportunistic scheduling for one type of ad-hoc networks where the wireless spectrum can be divided into multiple independent sub-channels for better efficiency. We started with a naive multi-channel protocol where the scheduling scheme is working independently from sub-channel to sub-channel. We showed that the naive protocol can only marginally improve the system throughput. We then developed a protocol to jointly consider the opportunistic scheduling behavior across multiple sub-channels. We characterized the optimal stopping rule and presented several bounds for the system throughput of the multi-channel protocol. We showed that by joint optimization of the scheduling scheme across multiple sub-channels, the proposed protocol improves the system throughput considerably in contrast to that of single-channel systems.

Performance aware cross-layer design in wireless multi hop networks. Further pursuing these novel ideas and methods, during this reporting period, we studied, analyzed and evaluated a performance-aware cross-layer design approach for wireless multihop networks. Through Network Utility Maximization (NUM) and weighted network graph modeling, a cross-layer algorithm for performing jointly routing, scheduling and congestion control was introduced. The performance-awareness is achieved by both the appropriate definition of the link weights for the corresponding applications requirements and the introduction of a weighted backpressure routing/scheduling. Contrary to the conventional backpressure the proposed algorithm scales the congestion gradients with the appropriately defined link weights. We analytically proved the queue stability achieved by the proposed cross-layer scheme, while its convergence to a close neighborhood of the optimal source rates values was proved via a novel ϵ -subgradient approach. Finally, through modeling and simulation, we demonstrated the performance improvements that can be achieved by the pro-

posed approach - when compared against existing methodologies in the literature - for two different examples with diverse application requirements, emphasizing respectively on delay and trustworthiness.

A Finite Memory EncoderDecoder Scheme for Control Under Communication Constraints.

This work proposes a remarkably simple algorithm based on a finite memory encoder and decoder capable of asymptotically stabilizing a first order discrete time linear system when the measurements from the plant are collected via a discrete memoryless channel. Given a performance specification for the closed loop system, we will characterize a set of discrete memoryless channels for which our scheme can match such performances and describe how to design the involved parameters accordingly.

Modeling and Detecting Community Hierarchies. A central problem in data mining and social network analysis is determining communities (clusters) among individuals or objects in the absence of external identification or tagging. In this work we propose a theoretical model that explicitly formalizes both the tight connections within each community and the hierarchical nature of the communities. In our model, each member of a community falls into a sub-community, and the sub-communities within this community have active interactions with each other while entities outside this community have fewer interactions with members inside. Given this formalization, we then propose an efficient algorithm that detects all the communities in this model, and prove that all the communities form a hierarchy. Empirical evaluations demonstrate that our formalization successfully models real world communities, and our algorithm compares favorably with existing approaches.

2.5 Evolving and uncertain environments

Global games with noisy information sharing: It is known that global games with noisy sharing of information do not admit a certain type of threshold policies. Motivated by this result, we investigate the existence of threshold-type policies on global games with noisy sharing of information and show that such equilibrium strategies exist and are unique if the sharing of information happens over a sufficiently noisy environment. To show this result, we establish that if a threshold function is an equilibrium strategy, then it will be a solution to a fixed point equation. Then, we show that for a sufficiently noisy environment, the functional fixed point equation leads to a contraction mapping, and hence, its iterations converge to a unique continuous threshold policy.

Planning for decentralized control of multiple robots under uncertainty We developed a probabilistic framework for synthesizing control policies for general multi-robot systems that is based on decentralized partially observable Markov decision processes (Dec-POMDPs). Dec-POMDPs are a general model of decision-making where a team of agents must cooperate to optimize a shared objective in the presence of uncertainty. Dec-POMDPs also consider communication limitations, so execution is decentralized. While Dec-POMDPs are typically intractable to solve for real-world problems, we recently developed a representation that uses macro-actions. Macro-actions more naturally represent real-world behavior which may be complex and require multiple time-steps (e.g., navigation to a waypoint or waiting for another agent). With our macro-action

approach, we significantly increased the size of problem that can be practically solved. In contrast to most existing methods that are specialized to a particular problem class, our approach can synthesize control policies that exploit any opportunities for coordination that are present in the problem, while balancing uncertainty, sensor information, and information about other agents. We used three variants of a warehouse task to show that a single planner of this type can generate cooperative behavior using task allocation, direct communication, and signaling, as appropriate. We also developed more scalable algorithms that search over policies represented as finite-state controllers. The resulting policy search algorithms require a simple simulator that models only the outcomes of executing a given set of motor controllers, not the details of the executions themselves and can to solve significantly larger problems than existing methods. Overall, we demonstrated that our framework and approaches can automatically optimize control and communication policies for complex multi-robot systems.

Scalable learning methods for multi-agent decision-making under uncertainty We developed scalable learning methods for the common case when a full model of the environment and other agents is not available (e.g., interaction with an unknown object, new team members, unknown humans). We developed methods that learn controllers for each agent from a small amount of data (generated from a simulator or real-world experience) in large decentralized POMDPs (Dec-POMDPs). A Bayesian nonparametric approach is used to learn both the size of the controllers and the parameters which produce the highest value. We showed that the algorithm scales to large problems while outperforming other state-of-the-art methods. We also developed methods where agents learn while acting in noisy, partially observable domains. This method simulates possible futures to choose the action that will optimally balance gaining information about the model and optimizing the underlying problem. As more time is available to simulate future outcomes, the learning method will improve (converging to the optimal solution). Experimental results show that we are able to provide high quality solutions to large multi-agent learning problems.

New Notions of Stability of Evolutionary Games: In collaboration between Martins and Shamma, we are exploring new notions of solution of an evolutionary game for cases when certain passivity conditions hold. We have made some preliminary progress reported in [NM4], where we propose the new notions of integrability and non-positive correlation and show that they have computational advantages and are equivalent to passivity. This will allow us to study new types of population games, such as when there is a delay between the state and the corresponding payoff.

Convergence in matching markets. Online matching markets are becoming increasingly important for matching buyers and sellers. Standard methods for finding optimal matches require full information about the value of all potential matches, which in practice is not available. We show that there exist simple distributed learning algorithms that achieve near-optimality in polynomial time even when market participants have no information about others valuations. The proofs use novel techniques characterizing random walks on trees.

Convergence rates in weakly acyclic games. Previous research has shown that there exist fully distributed learning algorithms that converge rapidly to equilibrium for congestion games, but to date little is known about convergence rates in other classes of weakly acyclic games, such as

coordination games, super-modular games, and dominance solvable games. We show that weak acyclicity is not sufficient to guarantee fast convergence; in fact there exist very simple weakly acyclic games such that virtually any distributed payoff-based learning algorithm takes exponentially long (in the population size) to come close to Nash equilibrium. However if the system is subject to aggregate shocks the learning process can be greatly accelerated. The source of these shocks could be temporary blockages in communications so that agents cannot see the payoffs from alternative strategies, or they could represent stochastic switching costs that make it costly to switch between strategies. We derive a bound on the waiting time for the learning process to come ϵ -close to Nash equilibrium from an arbitrary initial state ; this bound is exponential in the number of strategies, polynomial in $1/\epsilon$ and independent of the population size. This paper is forthcoming in *Econometrica*.

Achieving system optimality with fully distributed learning. We demonstrate a simple payoff-based learning rule that is completely decentralized, and that leads to an efficient configuration of actions in any n -person finite strategic-form game with generic payoffs. The algorithm requires no communication: agents respond solely to changes in their own realized payoffs, which are affected by the actions of other agents in the system in ways that they do not necessarily understand. The method can be applied to the optimization of complex systems with many distributed components, such as the routing of information in networks and the design and control of wind farms. The proof of the proposed learning algorithm relies on the theory of large deviations for perturbed Markov chains.

Stochastic stability analysis in games with many players. We consider large populations of agents who adjust their behavior using stochastic learning rules. We analyze the behavior of the system when deviations from best reply have vanishingly small probability and the population size is large. Transitions between equilibria can be described as solutions to continuous optimal control problems. These are used in turn to characterize the asymptotics of the stationary distribution, and so to determine the stochastically stable states.

Robust Distributed Routing in Dynamical Networks with Cascading Failures. Strong resilience properties of dynamical flow networks are analyzed for distributed routing policies. The latter are characterized by the property that the way the inflow at a non-destination node gets split among its outgoing links is allowed to depend only on local information about the current particle densities on the outgoing links. The strong resilience of the network is defined as the infimum sum of link-wise flow capacity reductions under which the network cannot maintain the asymptotic total inflow to the destination node to be equal to the inflow at the origin. A class of distributed routing policies that are locally responsive to local information is shown to yield the maximum possible strong resilience under such local information constraints for an acyclic dynamical flow network with a single origin-destination pair. The maximal strong resilience achievable is shown to be equal to the minimum node residual capacity of the network. The latter depends on the limit flow of the unperturbed network and is defined as the minimum, among all the non-destination nodes, of the sum, over all the links outgoing from the node, of the differences between the maximum flow capacity and the limit flow of the unperturbed network. We propose a simple convex optimization problem to solve for equilibrium limit flows of the unperturbed network that minimize

average delay subject to strong resilience guarantees, and discuss the use of tolls to induce such an equilibrium limit flow in transportation networks. Finally, we present illustrative simulations to discuss the connection between cascaded failures and the resilience properties of the network.

Partially observable stochastic games. Partially observable stochastic games (POSGs) and decentralized partially observable Markov decision processes (Dec-POMDPs) are rich models for competitive and cooperative decision-making under uncertainty, but providing solutions for these models is often intractable. We studied ways the models can be decomposed by formalizing the interface between local models as the influence agents can exert on one another. The resulting influence-based abstraction substantially generalizes previous work on exploiting weakly-coupled agent interaction structures. The influence-based approaches that we generalize have shown promising improvements in the scalability of planning for more restrictive models. Thus, our theoretical result serves as the foundation for practical algorithms that we anticipate will bring similar improvements to more general planning contexts, and also into other domains such as approximate planning, decision-making in adversarial domains, and online learning. Another approach for solving large Dec-POMDPs is by including realistic assumptions. The transition and observation independent Dec-MDP is a general subclass that has been shown to have complexity in NP (rather than NEXP), but optimal algorithms for this subclass are still inefficient in practice. We provided an updated proof that an optimal policy does not depend on the histories of the agents, but only the local observations (i.e., local states). We also developed a new algorithm based on heuristic search that is able to expand search nodes by using constraint optimization. This approach can increase scalability in terms of number of agents and problem size.

Propagation in large networks. We have developed a new model that falls between the very coarse differential equation models for contagion that are based on a fully mixed assumption for a large population and the very detailed simulation models that explicitly include every agent (individual in a population) as a node in a network to be studied. The new model includes both a network graph and averaged, continuous subsystems representing communities within the network. Thus the model is one of networked communities, which is a reasonable generalization of the well-mixed model that allows for heterogeneity while still being simple enough to permit analysis. This is important since the ability to obtain analytical results can be very useful in determining cause-effect relationships and in designing control schemes. We have studied the dynamics of the new models including analysis of epidemic threshold effects and comparison to available percolation theory results for the steady state of the system. We are also studying new and innovative problems in consensus modeling, such as the influence of external information inputs to agents on consensus in a network of communicating agents, the design of monitoring and control systems for the detection and isolation of malicious agents in a consensus network, and development of aggregate trust indices for agents from the point of view of an overall network.

Dynamic modes in large uncertain dynamic networks. We have continued our investigation into a new approach to modal participation analysis of dynamic networks. In 2009 we published a paper that announced that there was a dichotomy in modal participation analysis, while previously mode in state participation was assumed to be equivalent to state in mode participation for a dynamic system. The dichotomy is being found to have important and fundamental implications for

a wide range of issues in uncertain dynamic networks, including sensor placement, contingency analysis, detection of oscillations, and controller placement.

Utility design for engineering systems. Optimizing the price of anarchy: Game theory is emerging as a popular tool for distributed control of multiagent systems. In order to take advantage of these game theoretic tools the interactions of the autonomous agents must be designed within a game theoretic environment. A central component of this game theoretic design is the assignment of a local objective function to each decision-maker. One promising approach to utility design is assigning each agent an objective function in accordance with the agents Shapley value. This method frequently results in games that possess many desirable features including the existence and efficiency of pure Nash equilibria. In this paper we explore the relationship between the Shapley value utility design and the resulting efficiency of both pure Nash equilibria and coarse correlated equilibria. To study this relationship we introduce a simple class of resource allocation problems. Within this class of resource allocation problems, we derive an explicit relationship between the structure of the resource allocation problem and the efficiency of the resulting equilibria. Lastly, we derive a bicriteria bound for this class of resource allocation problems. By bicriteria bound, we mean a bound on the value of the optimal allocation relative to the value of an equilibrium allocation with additional agents.

State based potential games. A new design paradigm for networked control systems: There is a growing interest in the application of game theoretic methods to the design and control of multiagent systems. However, the existing game theoretic framework possesses inherent limitations with regards to these new prescriptive challenges. In this paper we propose a new framework, termed state based potential games, which introduces an underlying state space into the framework of potential games. This state space provides a system designer with an additional degree of freedom to help coordinate group behavior and overcome these limitations. Within the context of state based potential games, we characterize the limiting behavior of two learning algorithms termed finite memory better reply processes and log-linear learning. Lastly, we demonstrate the applicability of state based potential games on two cooperative control problems pertaining to distributed resource allocation and the design of local and distributed control laws.

Game design for distributed optimization: The central goal in multiagent systems is to design local control laws for the individual agents to ensure that the emergent global behavior is desirable with respect to a given system level objective. Ideally, a system designer seeks to satisfy this goal while conditioning each agents control law on the least amount of information possible. Unfortunately, there are no existing methodologies for addressing this design challenge. The goal of this paper is to address this challenge using the field of game theory. Utilizing game theory for the design and control of multiagent systems requires two steps: (i) defining a local objective function for each decision maker and (ii) specifying a distributed learning algorithm to reach a desirable operating point. One of the core advantages of this game theoretic approach is that this two step process can be decoupled by utilizing specific classes of games. For example, if the designed objective functions result in a potential game then the system designer can utilize distributed learning algorithms for potential games to complete step (ii) of the design process. Unfortunately, designing agent objective functions to meet objectives such as locality of information and efficiency of

resulting equilibria within the framework of potential games is fundamentally challenging and in many case impossible. In this paper we develop a systematic methodology for meeting these objectives using a broader framework of games termed state based potential games (introduced above). State based potential games are an extension of potential games where an additional state variable is introduced into the game environment hence permitting more flexibility in our design space. Furthermore, state based potential games possess an underlying structure that can be exploited by distributed learning algorithms in a similar fashion to potential games hence providing a new baseline for our decomposition.

Maximizing system performance through distributed learning: In previous work we identified novel procedures for learning Nash equilibrium in general n -person games. These learning rules are fully decentralized in the sense that agents do not receive any input from a central controller. Furthermore they require minimal amounts of information about the behaviour of other agents in the system. The key idea is that agents have a state variable (content/discontent) that switches depending on the pattern of recent payoffs. Lower than expected payoffs trigger the discontent state which involves frequent search; high payoffs lead to the content state with high probability (in the content state search is infrequent). In recent work with Jason Marden and Lucy Pao we show that this class of fully decentralized learning rules leads to maximum welfare configurations of the system with no central controller. Potential applications include the maximization of power from wind farms.

Learning efficient equilibria in decentralized matching markets: Decentralized matching platforms on the internet allow large numbers of agents to interact anonymously at virtually no cost. Very little information is available to market participants and trade takes place at many different prices simultaneously. We propose a decentralized learning process in such environments that leads to stable and efficient outcomes. Agents on each side of the market make demands of potential partners and are matched if their demands are mutually compatible. Matched agents occasionally experiment with higher demands, while unmatched agents lower their demands in the hope of attracting partners. This learning process implements core allocations even though agents have no knowledge of other agents' strategies, payoffs, or the structure of the game, and there is no central authority with such knowledge either.

Speed of convergence of learning rules in large populations: Consider a large population of agents who update their behaviors based on noisy payoff information about alternative actions: those with higher payoffs are chosen more frequently than those with lower payoffs. The resulting dynamical system exhibits fast convergence if the expected waiting time until the process comes near a Nash equilibrium is bounded above for all sufficiently large population sizes. Using stochastic approximation theory we show that fast convergence holds for a number of important classes of games, including coordination games, potential games, and supermodular games.

Game couplings: Learning dynamics and applications: Modern engineering systems consist of multiple coupled subsystems. Such subsystems are designed with local (possibly conflicting) goals, with little or no knowledge of the implementation details of other subsystems. Despite the ubiquitous nature of such systems very little is formally known about their properties and

global dynamics. We investigate such distributed systems by introducing a novel game-theoretic construct, that we call game- coupling. Game coupling intuitively allows us to stitch together the payoff structures of sub- games. In this work, we examine when do game couplings preserve or enhance desirable properties of the original games, such as convergence of best response dynamics and low price of anarchy.

Learning valuation functions: A core element of microeconomics and game theory is that consumers have valuation functions over bundles of goods and that these valuation functions drive their purchases. In particular, the value given to a bundle need not be the sum of values on the individual items but rather can be a more complex function of how the items relate. Common valuation classes considered in the literature include OXS, submodular, XOS, and subadditive valuations. Typically it is assumed that these valuations are known to the center or come from a known distribution. In this work we initiate the study of the approximate learnability of valuation classes in a distributional learning setting. We prove upper and lower bounds on the approximation guarantees achievable for learning over general data distributions by using only a polynomial number of examples. We provide a nearly tight learning guarantee on the efficient learnability of subadditive valuations that satisfy the property that the value given to a bundle is at most the sum of values on the individual items. We also provide nearly tight bounds for the subclass of XOS (aka fractionally subadditive) valuations, also widely used in the Algorithmic Game Theory literature. We additionally leverage the structure of valuations in a number of interesting subclasses and obtain algorithms with stronger learning guarantees. As an example, we use the fact that the XOS valuations have an appealing syntactic description (namely an XOS function can be represented as a depth-two tree with a MAX root and SUM inner nodes, where each such SUM node has a subset of items with associated positive weights as leaves) and provide a target-dependent learnability result for XOS functions; in particular, we provide efficient algorithms with stronger guarantees for XOS functions of polynomial description length. To do so, we first prove a structural result showing that a XOS function can be approximated well by the L -th root of a degree- L polynomial over the natural feature representation of the set S , and then using this fact we reduce the problem of learning such functions to the problem of learning a linear separator in a new feature representation (algorithmically, one could use SVM with a polynomial kernel). Our work combines central issues in economics with central issues in discrete convex optimization with central issues in learning theory.

Iterative auction design for graphical valuations: In multi-item iterative auctions, the auctioneer sets prices for the items she is selling, collects demands of the bidders at the current prices, and updates the prices in response to bidders demand reports. Such auctions are commonly used in practice, (for instance for selling spectrum, electricity, and natural resources) often without efficiency guarantees. On the other hand, the auction formats that are provably efficient (i.e., that implement an outcome which maximizes the social welfare) are either restricted to limited environments (where bidders view different items as substitutes), or they lead to complex pricing structures (e.g., the auctioneer may need to use a different price for each bundle). In this work, we focus on the question of designing simple efficient iterative auctions for multi-item settings that exhibit both value complementarity and substitutability. We start by restricting attention to a special class of valuation functions, which we refer to as graphical valuations. Value functions that

belong to this class are associated with a value graph, nodes of which correspond to the items that are being sold by the auctioneer. The edges of this graph capture the value complementarity and substitutability exhibited by the items. We first study the case where the underlying value graph is a tree, and provide an efficient iterative auction with a simple pricing structure. In particular, in the iterative auction we obtain, the auctioneer offers a single nonanonymous price for each item, and increases prices for the overdemanded items, and decreases them for underdemanded items until an efficient allocation is found. We then consider the case of more general graphical value functions, and show that for this case the efficient allocation problem can be formulated as the shortest path problem on an appropriately defined graph. Using the linear programming formulation of this shortest path problem, we obtain an iterative auction for general graphical valuations. The price structure used in the corresponding auction is simple when the underlying graph is similar to a path graph, and it becomes more complicated for graphs similar to complete graphs.

Power allocation policy for distributed estimation in wireless networks: An important problem in sensor networks is developing (distributed) algorithms for state estimation. When considering wireless sensor networks with limited energy available for communication, another important problem is designing power allocation schemes aimed at ensuring good estimation performance and network longevity. The problem increases in complexity if we impose the power allocation scheme to be obtained in a distributed manner. We have designed a distributed algorithm which computes a power allocation scheme aimed at ensuring accurate state estimates and at saving communication energy. The distributed algorithm resulted from solving an optimization problem with constraints. We analyzed how the power allocation scheme can be used for performing distributed estimation under different assumptions on the information shared between sensors.

Performance evaluation of the consensus-based distributed subgradient method: An important application of consensus algorithms, is to use them as a component in a distributed optimization scheme. We investigated collaborative optimization of an objective function expressed as a sum of local convex functions, when the agents make decisions in a distributed manner using local information, while the communication topology used to exchange messages and information is modeled by a graph-valued random process, assumed independent and identically distributed. Specifically, we studied the performance of the consensus-based multi-agent distributed subgradient method and showed how it depends on the probability distribution of the random graph. For the case of a constant stepsize, we first obtained an upper bound on the difference between the objective function, evaluated at the agents estimates of the optimal decision vector, and the optimal value. Second, for a particular class of convex functions, we obtained an upper bound on the distances between the agents estimates of the optimal decision vector and the minimizer. In addition, we obtained an expression for the rate of convergence to zero of the time varying component of the aforementioned upper bound. The addressed metrics were evaluated via their expected values. As an application, we showed how the distributed optimization algorithm can be used to perform collaborative system identification and provided numerical experiments under the randomized and broadcast gossip protocols.

Learning behaviors for coordination in stochastic networked systems: We considered the problem of coordination for efficient joint decision making in networks of autonomous agents.

When making a decision, an agent is influenced by its knowledge about others behaviors. Agents understanding of others behaviors is shaped through observing their actions over a long time. We have modeled the decision making on whether to cooperate in a group effort as a result of a series of two-person games between agents, where the payoff of each agent is computed as the sum of its payoffs from each of these games. The agents initially have different behaviors. In order to maximize their pay-off, they need to learn the others behavior and coordinate with them. We considered a behavior learning algorithm for a class of behavior functions and studied its effects on the emergence of coordination in the network. The conditions under which the learning algorithm converges were studied. We showed that for a class of linear functions the learning algorithm results in an extension of non-homogeneous consensus protocol to the more general case of block-stochastic matrices.

2.6 Networked operations

Norm-optimal control under sparse feedback interconnection: The sparsity pattern of a control system, which acts on a collection of subsystems, specifies the network of interconnections among local controllers (each mounted at a subsystem). The design of norm-optimal controllers constrained to a pre-specified sparsity pattern is a generalization of the long-standing open problem of designing optimal block-diagonal (decentralized) controllers. In 2010—2011, we proposed the first extension of Youlas classical parameterization so as to accommodate quadratically invariant (partially nested) sparsity constraints on the controller. It relies on the existence of a special factorization that leads to a convex parameterization of all sparsity-constrained stabilizing controllers. However desirable, the existence of such a special factorization cannot be predetermined in general. Moreover, there are counterexamples that show that a convex parameterization of all sparsity-constrained controllers may exist even when the aforementioned factorization does not. In 2011-2012, we obtained a significant extension of our prior results leading to a convex parameterization that does not rely on the aforementioned factorization. In addition, we have shown that the existence of a sparsity-constrained stabilizing controller is equivalent to the feasibility of an associated exact model-matching problem that is tractable by existing techniques. This result leads to the first systematic necessary and sufficient condition for the existence of a stabilizing controller satisfying a pre-specified quadratically invariant (partially nested) sparsity constraint. We have also proposed a systematic method to obtain a Youla-like parameterization of all stabilizing sparsity-constrained controllers that can be used to cast norm-optimal design problems as tractable convex programs. Unlike prior results, our method does not require knowledge of an “initial” stabilizing solution.

Design of LTI distributed observers: Here we consider a collection of LTI observers that are connected according to a pre-specified graph, and where each has access to a sub-partition of the output of an LTI plant. The objective is to design observers that achieve asymptotic omniscience of state of the plant, i.e., the estimate at each observer must converge to the state of the plant. We provide necessary and sufficient conditions for the existence of such observers. Our method is constructive and is rooted on the analysis of the fixed modes of an associated feedback system. Unlike existing results, the fact that our distributed observers are LTI facilitates the use of frequency domain methods for performance analysis in the presence of noise.

Convergence results for the agreement problem on Markovian random topologies: We investigated the linear distributed asymptotic agreement (consensus) problem for a network of dynamic agents whose communication network is modeled by a randomly switching graph. The switching is determined by a finite state, Markov process, each topology corresponding to a state of the process. We addressed both the cases where the dynamics of the agents is expressed in continuous and discrete time. We showed that, if the consensus matrices are doubly stochastic, convergence to average consensus is achieved in the mean square and almost sure sense, if and only if the graph resulting from the union of graphs corresponding to the states of the Markov process is strongly connected. The aim of our work is to show how techniques from the theory of Markovian jump linear systems, in conjunction with results inspired by matrix and graph theory, can be used to prove convergence results for stochastic consensus problems.

Consensus-based linear distributed filtering: We investigated the consensus-based distributed linear filtering problem, where a discrete time, linear stochastic process is observed by a network of sensors. We assumed that the consensus weights are known and we first provided sufficient conditions under which the stochastic process is detectable, i.e. for a specific choice of consensus weights there exists a set of filtering gains such that the dynamics of the estimation errors (without noise) is asymptotically stable. Next, we developed a distributed, sub-optimal filtering scheme based on minimizing an upper bound on a quadratic filtering cost. In the stationary case, we provided sufficient conditions under which this scheme converges; conditions expressed in terms of the convergence properties of a set of coupled Riccati equations. Further we introduced a consensus-based distributed filter, executed by a sensor network, inspired by the Markovian jump linear system filtering theory. We showed that the optimal filtering gains of the Markovian jump linear system can be used as an approximate solution of the optimal distributed filtering problem. This parallel allowed us to interpret each filtering gain corresponding to a mode of operation of the Markovian jump linear system as a filtering gain corresponding to a sensor in the network. The approximate solution can be implemented distributively and guarantees a quantifiable level of performance.

Asymptotic consensus problem on convex metric spaces: We introduced and investigated a significant generalization of the consensus problem. A consensus problem consists of a group of dynamic agents who seek to agree upon certain quantities of interest. The agents exchange information according to a communication network modeled as a directed time-varying graph. A convex metric space is a metric space on which we define a convex structure. In our work we generalized the asymptotic consensus problem to convex metric spaces. Under weak connectivity assumptions, we showed that if at each iteration an agent updates its state by choosing a point from a particular subset of the generalized convex hull generated by the agents current state and the states of its neighbors, then agreement is achieved asymptotically. As an application example, we introduced a probabilistic algorithm for reaching consensus on discrete finite sets and showed that it in fact fits our general framework. We further considered a variety of probabilistic versions of this extension.

Consensus problem under vanishing communication weights: Another aspect of typical work on distributed consensus problems is the neglect of the effects of vanishing communications be-

tween the agents. We revisited the classical multi-agent distributed consensus problem under the dropping of the assumption that the existence of a connection between agent implies weights uniformly bounded away from zero. We formulated and studied the problem by establishing global convergence results in discrete time, under fixed, switching and random topologies. We studied the application of the results to flocking networks.

Network formation and learning efficient topologies in collaborative control: We studied the collaboration of autonomous agents of mobile autonomous network subject to constraints, such as communication range constraints and power consumption constraints. Networks of mobile agents (e.g. autonomous vehicles) rely heavily on wireless communications as well as sensing devices for distributed path planning and decision making. Designing energy efficient distributed decision making algorithms in these systems is challenging and requires that different task-oriented information becomes available to the corresponding agents in a timely and reliable manner. We developed a systems engineering oriented approach to the design of networks of mobile autonomous systems, in which a cross-layer design methodology determines what structures are to be used to satisfy different task requirements. We identified a three-tier organization of these networks consisting of connectivity, communication, and action graphs and studied the interaction between them. It is expected that in each functionality of a network, there are certain topologies that facilitate better achievement of the agents objectives. Inspired from biological complex networks, we proposed a bottom-up approach in network formation, in which small efficient sub-graphs (motifs) for different functionalities of the network are determined. A simulation test bed was used to learn efficient sub-graphs (motifs) for local level communication. A dictionary of valid motifs was provided and an algorithm for derivation of task dependent motifs was implemented. The group mission consists of several tasks that the agents participate in. After the motif learning/extracting phase, a hierarchical method of integrating the sub-graphs into efficient global topologies was implemented. The overall network is thus formed as a combination of these sub-graphs. We showed that the bottom-up approach to network formation is capable of generating efficient topologies for multi-tasking complex environments.

Communication network challenges for collaborative vehicles: Networked systems of autonomous mobile agents have emerged in a variety of applications such as collaborative robotics, unmanned aerial/ground vehicles, mobile sensor networks and disaster relief operations. These agents utilize wireless communications for distributed computing, control and decision-making. Due to their limitations on energy supply, design and implementation of efficient distributed algorithms are crucial for these systems. We reviewed different design aspects of networks of collaborative vehicles. We surveyed existing methods of addressing physical (PHY) and medium access control (MAC) layer, routing, and geometric connectivity issues for such networks and showed limitations of these designs that address only single layer issues. We argued that a system engineering framework is needed for the design of such systems.

Expander families as information patterns for distributed control of vehicle platoons: Part of our research is focused on identifying best topologies for communication in multi-agent multi-tasking systems, including learning such topologies. We revisited the classical, and recently intensively studied, problem of the distributed control of a 1-D vehicle platoon. A state space for-

mulation was developed with the main objective being to understand the effects of the underlying information exchange pattern, between the vehicles, on the control performance of the platoon. The symmetric control case was considered where each vehicle gives equal weight to all the information available to it in determining its control law. We showed that expander families of graphs when used as information patterns result in stability margins decaying to zero at rate of at most $\mathcal{O}(1/N)$; an improvement over the previously known $\mathcal{O}(1/N^2)$ decay with nearest neighbor type information patterns.

Distributed opportunistic scheduling protocol for multi-channel wireless ad-hoc networks:

In heavily constrained communication environments in multi-agent systems, agents must learn to identify opportunities for communication and patterns of such opportunities. This leads to opportunistic (real-time) scheduling problems for communication between the agents. The topic of opportunistic scheduling for wireless ad-hoc networks has been studied for single-channel networks in several recent works. Since now many wireless systems provide multiple channels for data transmission, the opportunistic scheduling problem is of practical interest for networks with multiple channels. We studied the problem of opportunistic scheduling for one type of ad-hoc networks where the wireless spectrum can be divided into multiple independent sub-channels for better efficiency. We started with a naive multi-channel protocol where the scheduling scheme is working independently from sub-channel to sub-channel. We showed that the naive protocol can only marginally improve the system throughput. We then developed a protocol to jointly consider the opportunistic scheduling behavior across multiple sub-channels. We characterized the optimal stopping rule and presented several bounds for the system throughput of the multi-channel protocol. We showed that by joint optimization of the scheduling scheme across multiple sub-channels, the proposed protocol improves the system throughput considerably in contrast to that of single-channel systems.

Throughput-delay trade-off in wireless multi-hop networks via greedy hyperbolic embedding:

In multi-agent systems the communication network (and consequently its topology) should be robust and designed with several metrics in mind. For example some tasks require excellent delay performance, others require excellent throughput performance. Other metrics of interest include energy efficiency, trust and several others. Thus it is important to develop methodologies capable of analyzing, and even controlling, the tradeoffs between the various metrics. Real time applications over wireless multihop networks, demand routing/scheduling algorithms that achieve desirable delay-throughput trade-offs, with high throughput and low end-to-end delay. The backpressure algorithm, has received much attention by the research community in the past few years, as it satisfies the throughput optimal requirement. The backpressure algorithm performs routing and scheduling based on congestion gradients, by allowing transmission to the links that maximize the differential backlog between neighboring nodes. However, by deploying routing without using any information about the position or distance to the destination, it explores all possible source-destination paths leading to undesirable high delays. In our recent work, we proposed a new method of restricting the number of paths used by the backpressure algorithm, with the aid of the greedy embedding of a network in the hyperbolic space. We proposed two algorithms, the ‘‘Greedy’’ backpressure and the Greediest backpressure for both static and dynamic networks, which consider the network embedded in the hyperbolic space and combine the greedy routing in

hyperbolic coordinates with the backpressure scheduling. We proved analytically their throughput optimality and studied through simulations the induced improvement in the delay-throughput trade-off compared with the backpressure algorithm. We also investigated the coveted throughput-delay trade-off in static wireless multihop networks based on a “computer-aided” design of the backpressure scheduling/routing algorithm for networks embedded in hyperbolic space. Both routing and scheduling exploit the hyperbolic distances to orient the packets to the destination and prioritize the transmissions correspondingly. The proposed design provides us with the freedom of controlling its theoretical throughput optimality and of counterbalancing its practical performance through simulations, leading to significant improvements of the throughput-delay trade-off.

2.7 Evolving and uncertain environments

Robust distributed routing in dynamical networks: We consider dynamical formulation of network flows, where the network is modeled by a system of ordinary differential equations derived from mass conservation laws on directed acyclic graphs with a single origin-destination pair and a constant inflow at the origin. The rate of change of the particle density on each link of the network equals the difference between the inflow and the outflow on that link. The inflow to a link is determined by the total flow arriving to the tail node of that link and the routing policy at that tail node. The outflow from a link is modeled to depend on the current particle density on that link through a flow function. Every link is assumed to have finite capacity for particle density and the flow function is modeled to be strictly increasing as density increases from zero up to a certain point and is then strictly decreasing until it becomes zero for the maximum particle density. We derived new results on resilience of such networks under distributed routing policies towards perturbations that reduce link-wise flow functions. In particular, we propose an algorithm to compute an upper bound on the maximum resilience over all distributed routing policies, and a lower bound based on the analysis of a specific routing policy. These results highlight the role of cascades on the resilience of the network.

Clustering stable instances: Problems of clustering data from pairwise distance information are ubiquitous in science. A common approach for solving such problems is to view the data points as nodes in a weighted graph (with the weights based on the given pairwise information), and then to design algorithms to optimize various objective functions such as k -median or min-sum. For example, in the k -median clustering problem the goal is to partition the data into k clusters C_i , giving each a center c_i , in order to minimize the sum of the distances of all data points to the centers of their cluster. In the min-sum clustering approach the goal is to find k clusters C_i that minimize the sum of all intra-cluster pairwise distances. Yet unfortunately, for most natural clustering objectives, finding the optimal solution to the objective function is NP-hard. As a consequence, there has been substantial work on approximation algorithms with both upper and lower bounds on the approximability of these objective functions on worst case instances. Distances between data points in clustering instances are often based on a heuristic measure, and accordingly, interesting instances should be resilient to small perturbations in these distances. In particular, if small perturbations can cause the optimum clustering for a given objective to change drastically, then that probably is not a meaningful objective to be optimizing. Motivated by such considerations, recent work proposed analyzing objective based clustering problems under the assumption that the opti-

minimizing clustering to the objective is preserved under small multiplicative perturbations to distances between points. In our work, we provide several results within this framework. For separable center-based objectives (which include the classic k -means and k -median objectives), we present an algorithm that can optimally cluster instances resilient to $1 + \sqrt{2}$ -factor perturbations, solving an open problem. For the k -median objective, we additionally give algorithms for a weaker, relaxed, and more realistic assumption in which we allow the optimal solution to change in a small fraction of the points after perturbation. We also provide positive results for min-sum clustering which is a generally much harder objective than k -median (and also non-center-based). Our algorithms are based on new linkage criteria that may be of independent interest.

Learning in Markov games: Partially observable stochastic games (POSGs) and decentralized partially observable Markov decision processes (Dec-POMDPs) are rich models for competitive and cooperative decision-making under uncertainty, but providing solutions for these models is often intractable. We studied ways the models can be decomposed by formalizing the interface between local models as the influence agents can exert on one another. The resulting influence-based abstraction substantially generalizes previous work on exploiting weakly-coupled agent interaction structures. The influence-based approaches that we generalize have shown promising improvements in the scalability of planning for more restrictive models. Thus, our theoretical result serves as the foundation for practical algorithms that we anticipate will bring similar improvements to more general planning contexts, and also into other domains such as approximate planning, decision-making in adversarial domains, and online learning. Another approach for solving large Dec-POMDPs is by including realistic assumptions. The transition and observation independent Dec-MDP is a general subclass that has been shown to have complexity in NP (rather than NEXP), but optimal algorithms for this subclass are still inefficient in practice. We provided an updated proof that an optimal policy does not depend on the histories of the agents, but only the local observations (i.e., local states). We also developed a new algorithm based on heuristic search that is able to expand search nodes by using constraint optimization. This approach can increase scalability in terms of number of agents and problem size.

Trust-based multi-agent filtering for increased smart grid security: We addressed the problem of state estimation of the power system for the Smart Grid. We assumed that the monitoring of the electrical grid is done by a network of agents with both computing and communication capabilities. We proposed a security mechanism aimed at protecting the state estimation process against false data injections originating from faulty equipment or cyber-attacks. Our approach is based on a multi-agent filtering scheme, where in addition to taking measurements, the agents are also computing local estimates based on their own measurements and on the estimates of the neighboring agents. We combined the multi-agent filtering scheme with a trust-based mechanism under which each agent associates a trust metric to each of its neighbors. These trust metrics are taken into account in the filtering scheme so that information transmitted from agents with low trust is disregarded. In addition, a mechanism for the trust metric update is also introduced, which ensures that agents that diverge considerably from their expected behavior have their trust values lowered. This problem represents an interesting learning problem as trust values are learned from observing the agents actions.

Selfish response to epidemic propagation: Spreading information via various means in a system of networked agents has substantial effects on their decision making. Learning is an important aspect of such problems. An epidemic that spreads in a network calls for a decision on the part of the network users. They have to decide whether to protect themselves or not. Their decision depends on the trade-off between the perceived infection and the protection cost. Aiming to help users reach an informed decision, security advisories provide periodic information about the infection level in the network. We studied the best-response dynamic in a network whose users repeatedly activate or de-activate security, depending on what they learn about the infection level. Our main result is the counterintuitive fact that the equilibrium level of infection increases as the users learning rate increases. The same is true when the users follow smooth best-response dynamics, or any other continuous response function that implies higher probability of protection when learning a higher level of infection. In both cases, we characterized the stability and the domains of attraction of the equilibrium points. Our finding also holds when the epidemic propagation is simulated on human contact traces, both when all users are of the same best-response behavior type and when they are of two distinct behavior types.

Reachability analysis for linear set-dynamics driven by random sets: An important challenge in the multi-layer framework for studying multi-agent systems is the ability to estimate what can be achieved under various constraints. Such problems can be formulated as reachability problems in appropriately chosen state spaces. The uncertainties involved lead to consideration of set-valued dynamical systems. We initiated such studies in where we studied linear setdynamics driven by random convex compact sets (RCCSs), and derived the setdynamics of the expectations of the associated reach sets as well as the dynamics of the corresponding covariance functions. We established that the expectations of the reach sets evolve according to deterministic linear setdynamics while the associated dynamics of covariance functions evolves on the Banach space of continuous functions on the dual unit ball. The general framework was specialized to the case of Gaussian RCCSs, and it was shown that the Gaussian structure of random sets is preserved under linear setdynamics of random sets.

Social networks over wireless networks: We considered the formation, operation and maintenance of dynamic social networks (among human users) supported by technological communication networks such as wireless networks, or hybrid wireline-wireless networks. The technological (physical) networks of interest display dynamic behavior in several dimensions, including variable connectivity, variable congestion, variable link characteristics. As broadband wireless devices and networks are becoming ubiquitous these human-machine systems, that combine the social aspects and behavioral activities of humans with the technological characteristics of the underlying physical networks, provide several important challenges in efforts to model them, evaluate their performance and dynamically control them so certain performance requirements are met. These include combinations of performance, trust, privacy, energy efficiency. We developed novel models for these complex human-machine systems that incorporate social network behavioral models and wireless network models that are inspired from statistical physics (hyperbolic graphs). We investigated the performance of wireless network protocols that support and respond to the constraints implied by the social network they support.

3 Acknowledgement/Disclaimer

This work was sponsored (in part) by the Air Force Office of Scientific Research, USAF, under grant/contract FA9550–09–1–0538. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

4 Personnel Supported During Duration of Grant

Note: *Italics* denotes past participant.

- Faculty:
 - E. Abed, U of Maryland College Park.
 - M.F. Balcan, Georgia Tech.
 - J.S. Baras, U of Maryland College Park.
 - M.A. Dahleh, MIT.
 - E.M. Feron, Georgia Tech.
 - L. Kaelbling, MIT.
 - J. Marden, U of Colorado Boulder.
 - N. Martins, U of Maryland College Park.
 - A. Ozdaglar, MIT.
 - J.S. Shamma, Georgia Tech.
 - H.P. Young, Johns Hopkins University & University of Oxford.
- Postdocs:
 - Christopher Amato, MIT.
 - Yahok Babichenko, Johns Hopkins University.
 - Pedram Hovareshti, U of Maryland College Park.
 - Kwang-Ki Kim, Georgia Tech.
 - Georgios Kotsalis, Georgia Tech.
 - Ion Matei, U of Maryland College Park.
 - Ruth Mehta, Georgia Tech.
 - Georgios Piliouras, Georgia Tech.
 - Christina Rose-Kryrtos, U of Maryland College Park.
 - Ruth Umer, Georgia Tech.

- Shah-An Yang, U of Maryland College Park.
 - *Itai Arieli, Johns Hopkins University. (current: Technion, Israel)*
 - *Florin Constantin, Georgia Tech. (current: A9.com)*
 - *Tao Jiang, University of Maryland College Park. (current: Research Scientist at Intel-ligent Automation, Inc.)*
 - *Heinrich H. Nax, Johns Hopkins University. (current: ETH, Zurich)*
 - *Frans Oliehoek, MIT.*
 - *Mathias Staudigl, Johns Hopkins University. (current: University of Bielefeld)*
- Graduate students:
 - Elie Adam, MIT.
 - Prateek Bhakta, Georgia Tech.
 - Hua Chen, U of Maryland College Park (to join Qualcomm Research and Develop-ment)
 - Chris Berlind, Georgia Tech.
 - Shang Tse Chen, Georgia Tech.
 - Bo Dai, Georgia Tech.
 - Steven Ehrlich, Georgia Tech.
 - Peixin Gao, U of Maryland College Park.
 - Guarang Gava, Georgia Tech.
 - Yonatan Gefen, U of Maryland College Park.
 - Jenny Jeong, Georgia Tech.
 - Yingyu Liang, Georgia Tech.
 - Yusun Lim, Georgia Tech. (to join Samsung)
 - Van Sy Mai, U of Maryland College Park.
 - Dipankar Maiti, U of Maryland College Park.
 - Ali Makhdoumi, MIT.
 - Anup Menon, U of Maryland College Park (to join General Electric Global Research)
 - Jacob Moschler, U of Maryland College Park.
 - Shinkyu Park, U of Maryland College Park.
 - Kaushik Patnaik, Georgia Tech.
 - Daniel Pickem, Georgia Tech.
 - *Gabriel Kreindler, Johns Hopkins University.*
 - Bary S.R. Pradelski, Postdoc, University of Oxford (Current: McKinsey & Co.).
 - Christina Rose-Kyrtos, U of Maryland College Park.

- Christoforos Somarakis, U of Maryland College Park.
 - Eleni Stai, U of Maryland College Park.
 - Doohyun Sung, U of Maryland College Park. (to join Qualcomm Research and Development)
 - Tuan Ta, U of Maryland College Park.
 - Paul Tschirhart, U of Maryland College Park.
 - Marcos Vasconcelos, U of Maryland College Park.
 - David Ward, U of Maryland College Park.
 - E. Wei, MIT.
 - Xiaoya Wei, U of Maryland College Park.
 - Miguel Venegaz Zambrana, U of Maryland College Park.
 - *Ozan Candogan, MIT. (Current: Assistant Professor at Duke University)*
 - *Pedram Hovareshti, U of Maryland College Park. (Current: Postdoctoral Fellow at U Maryland)*
 - *Ion Matei, U of Maryland College Park. (PhD 2010. Current: Postdoctoral Fellow at NIST & U Maryland)*
 - *Edward MacDonald, Georgia Tech. (MS 2011)*
 - *Serban Sabau, U of Maryland College Park (PhD 2011. Current: Assistant Professor at Stevens Institute of Technology).*
- Visitors:
 - Sasa Rakovic (worked and collaborated on the project, but not supported by project funds)
 - Undergraduate students:
 - Jamison Bryant, U of Maryland College Park.
 - Aidan Dowdle, Georgia Tech (volunteer researcher)
 - Nocilas Homble, U of Maryland College Park.
 - Austin Hou, U of Maryland College Park.
 - Sarpesh Hebbar (volunteer researcher)
 - Heather Zion (volunteer researcher)

5 Special Events

1. Culminating CDC workshop on Human-Machine Networks and Distributed Autonomy, December 2015.

6 Publications

1. E.H. Abed and B. Hamzi, “On participation factors for nonlinear systems”, *SIAM Conference on Control and Its Applications*, San Diego, July 2013.
2. D. Acemoglu, O. Candogan, A. Ozdaglar, and P.A. Parrilo, “Iterative Auction Design for Graphical Valuations”, working paper, 2012.
3. E. Adam, M.A. Dahleh,, A. Ozdaglar, “Towards an Algebra for Cascade Effects”, working paper, 2014.
4. I. Al-Shyoukh and J.S. Shamma, A coevolutionary model of strategic network formation, in P. Contucci et al. (eds):*Complex Networks V*, Springer, 2014.
5. A. Agha-mohammadi, S. Omidshafiei, C. Amato and J. P. How, “Graph-Based Planning to Solve Multi-Agent POMDPs”, *Proceedings of the Workshop on Distributed Control and Estimation for Robotic Vehicle Networks* at Robotics: Science and Systems Conference, July 2014.
6. C. Amato. Cooperative Decision Making. In *Decision Making Under Uncertainty: Theory and Application*, edited by Mykel J. Kochenderfer. MIT Press, Cambridge, MA, July, 2015.
7. C. Amato, “Cooperative Decision Making”, to appear in *Decision Making Under Uncertainty: Theory and Application*, edited by Mykel J. Kochenderfer. MIT Press, 2014.
8. C. Amato, G. D. Konidaris and L. P. Kaelbling, “Planning with Macro-Actions in Decentralized POMDPs”, In *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-14)*, May 2014.
9. C. Amato, G. Konidaris, J. P. How and L. P. Kaelbling, “Combined Planning Under Uncertainty for Communication and Control in Multi-Robot Teams”, *Proceedings of the Workshop on Communication-aware Robotics: New Tools for Multi-Robot Networks, Autonomous Vehicles, and Localization (CarNet)* at Robotics: Science and Systems Conference, July 2014.
10. C. Amato, G. D. Konidaris, G. Cruz, C. Maynor, J. P. How, L. P. Kaelbling, “Planning for Decentralized Control of Multiple Robots Under Uncertainty”, *Proceedings of the Workshop on Planning and Robotics (PlanRob)* at the Twenty-Fourth International Conference on Automated Planning and Scheduling (ICAPS-14), June 2014.
11. C. Amato, George Konidaris, Ariel Anders, Gabriel Cruz, Jonathan P. How and Leslie P. Kaelbling. Policy Search for Multi-Robot Coordination under Uncertainty. In *Proceedings of the 2015 Robotics: Science and Systems Conference (RSS-15)*, July 2015.
12. C. Amato, George D. Konidaris, Gabriel Cruz, Christopher Maynor, Jonathan P. How and Leslie P. Kaelbling. Planning for Decentralized Control of Multiple Robots Under Uncertainty. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA-15)*, May 2015.

13. C. Amato and Frans A. Oliehoek. Scalable Planning and Learning for Multiagent POMDPs. In *Proceedings of the Twenty-Ninth AAAI Conference on Artificial Intelligence (AAAI-15)*, January 2015.
14. I. Arieli and Y. Babichenko, “Average testing and the efficient boundary”, *Journal of Economic Theory* (to appear).
15. I. Arieli and H. P. Young, “Fast convergence in population games”, Discussion Paper 570, Department of Economics, University of Oxford (submitted for journal publication).
16. I. Arieli and M. Mueller-Frank, “Generic outcomes of observational learning”, Working Paper, University of Oxford, 2012.
17. I. Arieli and M. Mueller-Frank, “Inferring beliefs from actions”, working paper, University of Oxford, 2012.
18. I. Arieli and H.P. Young, “Stochastic learning dynamics in population games”, submitted to *Econometrica*, 2013.
19. I. Arieli and H. P. Young, “Stochastic Learning Dynamics and Speed of Convergence in Population Games”, working paper, 2014.
20. I. Arieli and H. P. Young, Stochastic Learning Dynamics and Speed of Convergence in Population Games, forthcoming in *Econometrica*.
21. P. Awasthi, M.-F. Balcan, A. Blum, O. Sheffet, and S. Vempala, “On Nash-equilibria of approximation-stable games”, *Current Science Journal*, vol. 103, 2012.
22. Y. Babichenko and L.J. Shulman, One-Shot Bargaining Mechanisms, Working Paper, Department of Industrial Engineering and Management, Technion, 2015.
23. M.-F. Balcan and P. Long, “Active and Passive Learning of Linear Separators under Log-concave Distributions”, *Proceedings of the 26th Annual Conference on Learning Theory (COLT)*, 2013.
24. M.-F. Balcan and Y. Liang, “Modeling and Detecting Community Hierarchies”, *Proceedings of the 2nd International Workshop on Similarity-Based Pattern Analysis and Recognition (SIMBAD)*, 2013.
25. M.-F. Balcan, C. Berlind, S. Ehrlich, and Y. Liang, “Efficient Semi-supervised and Active Learning of Disjunctions”, *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.
26. M.-F. Balcan, A. Blum, Y. Mansour, “Exploiting Ontology Structures and Unlabeled Data for Learning”, *Proceedings of the 30th International Conference on Machine Learning (ICML)*, 2013.
27. M.-F. Balcan, A. Blum, Y. Mansour, “Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior”, *SIAM J. Comput.*, vol. 42, 2013, pp. 230–264.

28. M.-F. Balcan, A. Blum, and Y. Mansour, “The price of uncertainty”, *ACM Transactions on Economics and Computation*, vol. 1, September 2013.
29. M.F. Balcan, F. Constantin, G. Piliouras, and J.S. Shamma, “Game-theoretic couplings and applications”, *IEEE Conference on Decision and Control*, Orlando, FL, 2011.
30. M.F. Balcan, F. Constantin, S. Iwata, and L. Wang, “Learning valuation functions”, *Conference on Learning Theory (COLT 2012)*, Edinburgh, Scotland, June 2012.
31. M.F. Balcan, E. Blais, A. Blum, and L. Yang, “Active Property Testing”, *53rd Annual Symposium on Foundations of Computer Science (FOCS 2012)*, New Brunswick, NJ, October 2012 (to appear).
32. M.F. Balcan, A. Blum, and Y. Mansour, “Circumventing the Price of Anarchy: Leading Dynamics to Good Behavior”, *SIAM Journal of Computing* (to appear).
33. M.F. Balcan and S. Hanneke, “Robust Interactive Learning”, *Conference on Learning Theory (COLT 2012)*, Edinburgh, Scotland, June 2012.
34. M.F. Balcan and Y. Liang, “Clustering under Perturbation Resilience”, *39th International Colloquium on Automata, Languages and Programming (ICALP 2012)*, Warwick, UK, July 2012.
35. M.-F. Balcan et al., “Learning Economic Parameters from Revealed Preferences”, *Conference on Web and Internet Economics*, December 2014. (to appear)
36. M.-F. Balcan et al., “Active learning and best-response dynamics”, *Neural Information Processing Systems Foundation*, December 2014. (to appear)
37. M.-F. Balcan et al., “Improved distributed principal component analysis”, *Neural Information Processing Systems Foundation*, December 2014. (to appear)
38. M.F. Balcan, A. Daniely, R. Mehta, Ruth Uner, V. V. Vazirani Learning Economic Parameters from Revealed Preferences. Conference on Web and Internet Economics (WINE), 2014
39. G. Baldan, M. Dahleh and A. Megretski, “A Finite Memory EncoderDecoder Scheme for Control Under Communication Constraints”, working paper, MIT, 2013.
40. J.S. Baras, P. Hovareshti and H. Chen, “Motif-based Communication Network Formation for Task Specific Collaboration in Complex Environments”, *American Control Conference (ACC 2011)*, San Francisco, CA, June/July 2011, pp. 1051–1056
41. J.S. Baras, “Network Tomography: New Rigorous Approaches for Discrete and Continuous Problems,” invited paper, *6th International Symposium on Communications, Control and Signal Processing (ISCCSP2014)*, Athens, Greece, May 21-23, 2014.
42. J.S. Baras, “A Fresh Look at Network Science: Interdependent Multigraphs Models Inspired from Statistical Physics,” invited paper, *6th International Symposium on Communications, Control and Signal Processing (ISCCSP2014)*, Athens, Greece, May 21-23, 2014.

43. A. Bellet et al., “Distributed Frank-Wolfe Algorithm: A Unified Framework for Communication-Efficient Sparse Learning”, *ICML 2014 workshop on New Learning Frameworks and Models for Big Data*, December 2014. (to appear)
44. A. Bellet, Y.Liang, A.B. Garakani, M.F. Balcan, F. Sha. A Distributed Frank-Wolfe Algorithm for Communication-Efficient Sparse Learning. *SIAM International Conference on Data Mining (SDM 2015)*.
45. M. Benaïm, W.H. Sandholm, and M. Staudigl, “Large deviations and stochastic stability in the large population limit”, working paper, University of Bielefeld, 2013.
46. H.P. Borowski, J.R. Marden, and J.S. Shamma, “Learning efficient coarse correlated equilibria”, *53rd IEEE Conference on Decision and Control*, Los Angeles, CA, December 2014. (to appear)
47. C. Berlind and R. Uner Active Nearest Neighbors in Changing Environments. *International Conference on Machine Learning (ICML)*, 2015.
48. H. Borowski, J.R. Marden, and E.W. Frew, “Fast Convergence in Semi-Anonymous Potential Games,” *IEEE Conference on Decision and Control*, 2013.
49. H. Borowski, J.R. Marden, D.S. Leslie and E.W. Frew, “Coarse Resistance Tree Methods For Stochastic Stability Analysis,” *IEEE Conference on Decision and Control*, 2013.
50. P.N. Brown and J.R. Marden, “Social Coordination in Unknown Price-Sensitive Populations,” *IEEE Conference on Decision and Control*, 2013.
51. O. Candogan, A. Ozdaglar, and P.A. Parrilo, “Iterative Auction Design for Graphical Valuations”, working paper, MIT, 2012.
52. O. Candogan, A. Ozdaglar, and P.A. Parrilo, “Near-Potential Games: Geometry and Dynamics”, *Transactions in Economics and Computations*, vol. 1, 2013.
53. O. Candogan, A. Ozdaglar, and P.A. Parrilo, “Dynamics in Near-Potential Games,” *Games and Economic Behavior*, vol. 82, pp. 65–90, 2013.
54. O. Candogan, A. Ozdaglar, and P.A. Parrilo, “Near-Potential Games: Geometry and Dynamics”, *Transactions in Economics and Computations* (to appear).
55. O. Candogan, A. Ozdaglar, and P.A. Parrilo, “Dynamics in Near-Potential Games,” *Games and Economic Behavior* (to appear).
56. O. Candogan, I. Menache, A. Ozdaglar, and P.A. Parrilo, “Flows and Decompositions of Games: Harmonic and Potential Games,” *Mathematics of Operations Research* (to appear).
57. G.C. Chasparis, J.S. Shamma, and A. Rantzer, Non-convergence to saddle boundary points under perturbed reinforcement learning, *International Journal of Game Theory*, September 2014 (online version).

58. G.C. Chasparis, A. Arapostathis, and J.S. Shamma, "Aspiration learning in coordination games", *SIAM Journal on Control and Optimization*, February 2013, pp. 465-490
59. G.C. Chasparis, A. Arapostathis, and J.S. Shamma, "Aspiration learning in coordination games", *SIAM Journal on Control and Optimization* (to appear).
60. G.C. Chasparis and J.S. Shamma, "Perturbed learning automata in potential games", *50th IEEE Conference on Decision and Control*, Orlando, FL, 2011.
61. G.C. Chasparis, A. Arapostathis, and J.S. Shamma, "Fair Scheduling in Common-Pool Games by Aspiration Learning", *8th International Workshop on Resource Allocation and Cooperation in Wireless Networks (RAWNET 2012)*, Paderborn, Germany, 2012.
62. H. Chen and J.S. Baras, "A Distributed Opportunistic Scheduling Protocol for Multi-Channel Wireless Ad-Hoc Networks," *Proceedings 2012 IEEE Globecom*, pp. 292-297, Anaheim, CA, December 3-7, 2012.
63. H. Chen and J.S. Baras, "Distributed Opportunistic Scheduling for Wireless Ad-Hoc Networks with Block-Fading Model," to appear in *IEEE Journal on Selected Areas in Communications*, 2013.
64. H. Chen, P. Hovareshti and J. S. Baras, "Communication Network Challenges for Collaborative Vehicles", invited paper, *Proceedings Asilomar Conference on Signals, Systems, and Computers*, Pacific Grove, CA, November 2011, pp. 1472-1476
65. H. Chen and J.S. Baras, "A Distributed Opportunistic Scheduling Protocol for Multi-Channel Wireless Ad-Hoc Networks," *2012 IEEE Globecom*, Anaheim, CA, December 2012 (to appear).
66. H. Chen and J.S. Baras, "Distributed Opportunistic Scheduling for Wireless Ad-Hoc Networks with Block-Fading Model," *IEEE Journal on Selected Areas in Communications*, Vol. 31, Issue: 11, pp. 2324-2337, November 2013.
67. G. Como, K. Savla, D. Acemoglu, M.A. Dahleh, and E. Frazzoli, "Robust Distributed Routing in Dynamical Networks with Cascading Failures", *IEEE Conference on Decision and Control*, Maui, HI, December 2012.
68. G. Como, K. Savla, D. Acemoglu, M.A. Dahleh and E. Frazzoli, "Stability Analysis of Transportation Networks with Multiscale Driver Decisions", (submitted for journal publication).
69. G. Como, K. Savla, D. Acemoglu, M.A. Dahleh and E. Frazzoli, "Robust Distributed Routing in Dynamical Networks - Part I: Locally Responsive Policies and Weak Resilience", *IEEE Transactions on Automatic Control*, Vol. 58, No. 2, 2013 (to appear).
70. G. Como, K. Savla, D. Acemoglu, M.A. Dahleh and E. Frazzoli, "Robust Distributed Routing in Dynamical Networks - Part II: Strong Resilience, Equilibrium Selection and Cascaded Failures", *IEEE Transactions on Automatic Control*, Vol. 58, No. 2, 2013 (to appear).

71. J.S. Dibangoye, Christopher Amato, Olivier Buffet and François Charpillet. Optimally Solving Dec-POMDPs as Continuous-State MDPs. To appear in the *Journal of Artificial Intelligence Research*, 2015.
72. J.S. Dibangoye, C. Amato and A. Doniec. “Scaling Up Decentralized MDPs Through Heuristic Search”. In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, August 2012.
73. B. Dai et al., “Scalable Kernel Methods via Doubly Stochastic Gradient”, *Neural Information Processing Systems Foundation*, December 2014. (to appear)
74. J. S. Dibangoye, C. Amato, O. Buffet and F. Charpillet, “Exploiting Separability in Multi-Agent Planning with Continuous-State MDPs”, *Proceedings of the Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-14)*, May 2014. [best paper award]
75. N. Du, Y. Liang, M.-F. Balcan, and L. Song, “Influence Function Learning in Information Diffusion Networks”, *International Conference on Machine Learning*, June 2014.
76. M.J. Fox and J.S. Shamma, “Language evolution in finite populations”, *50th IEEE Conference on Decision and Control*, Orlando, FL, 2011.
M.J. Fox, B. Touri, and J.S. Shamma, Dynamics in atomic signaling games, *Journal of Theoretical Biology*, July 2015, pp. 8290.
77. M.J. Fox and J.S. Shamma, Probabilistic performance guarantees for distributed self-assembly, *IEEE Transactions on Automatic Control*, December 2015.
78. M.J. Fox, G. Piliouras, and J.S. Shamma, “Medium and long-run properties of linguistic community evolution”, *9th International Conference on the Evolution of Language (Evolang IX)*, Kyoto, Japan, 2012.
79. Y. Gefen, *An Adaptively Sampled Path Planner Using Waypoints: An Any-Angle Variant*, MS Thesis, ECE Department, University of Maryland at College Park, 2014.
80. B. Gharesifard, B. Touri, T. Basar, and J.S. Shamma, On the convergence of piecewise linear strategic interaction dynamics on networks, *IEEE Transactions on Automatic Control*, 2015 (accepted).
81. R. Gopalakrishnan, J.R. Marden, and A. Wierman, “Potential Games are Necessary to Ensure Pure Nash Equilibria in Cost Sharing Games,” *Electronic Commerce*, 2013.
82. R. Gopalakrishnan, J.R. Marden, and A. Wierman, “Potential Games are Necessary to Ensure Pure Nash Equilibria in Cost Sharing Games,” submitted to *Mathematics of Operations Research*, 2013.
83. B. Hamzi and E.H. Abed, “Local mode-in-state participation factors for nonlinear systems,” *IEEE Conference on Decision and Control*, 2014 (to appear).

84. B. Hamzi and E.H. Abed, "Local analysis of modal participation in nonlinear systems," working paper, 2014.
85. B. Hamzi and E.H. Abed, "Local mode-in-state participation factors for nonlinear systems," 53rd IEEE Conference on Decision and Control, Los Angeles, December 15-17, 2014, pp.
86. B. Hamzi and E.H. Abed, "Local analysis of modal participation in nonlinear systems," Submitted to Automatica.
87. W.A. Hashlamoun and E.H. Abed, "Analysis and optimal control of nonlinear rumor propagation models", submitted for journal publication, 2013.
88. P. Hovareshti and J. S. Baras, Learning Behaviors for Coordination in Networked Stochastic Systems, submitted for journal publication, 2012.
89. P. Hovareshti and J. S. Baras, "Learning Behaviors for Coordination in Networked Systems and a Generalized Consensus Protocol," *IEEE Conference on Decision and Control*, Orlando, Florida, USA, December 2011, pp. 2247-2452.
90. P. Hovareshti and J. S. Baras, "Learning Behaviors for Coordination in Networked Stochastic Systems," working paper, 2013.
91. K.-K. Kim and J.S. Shamma, "Multiagent sequential hypothesis testing", *53rd IEEE Conference on Decision and Control*, Los Angeles, CA, December 2014. (to appear)
92. G. E. Kreindler and H. P. Young, "Rapid Innovation Diffusion in Social Networks," *Proceedings of the National Academy of Sciences*, 2014 vol 111 Suppl 3, 10881–10888.
93. G.E. Kreindler and H.P. Young, "Fast convergence in evolutionary equilibrium selection", *Games and Economic Behavior*, vol. 80, pp. 39–67, 2013.
94. G. Kotsalis, Limits of performance for the model reduction problem of Hidden Markov Models, *IEEE Conference on Decision and Control*, 2015.
95. C.R. Kyrtos and J.S. Baras, "Studying the Role of APOE in Alzheimers Disease Pathogenesis Using a Systems Biology Model," *Proceedings 5th International Conference on Bioinformatics and Computational Biology (BICoB2013)*, Honolulu, Hawaii, March 2013.
96. D. Leslie and J. R. Marden, "Equilibrium Selection in Potential Games with Noisy Pay-offs", *International Conference on NETwork Games, COntrol and OPTimization (NetG-Coop)*, Paris, France, October 2011.
97. N. Li and J.R. Marden, "Decoupling Coupled Constraints Through Utility Design", (submitted for journal publication).
98. N. Li and J.R. Marden, "Game Design for Distributed Optimization", *IEEE Conference on Decision and Control*, Orlando, FL, December 2011.
99. N. Li and J.R. Marden, "Game Design for Distributed Optimization with a Time Varying Communication Graph", *IEEE Conference on Decision and Control*, 2012 (to appear).

100. Y. Lim. *Game Theoretic Distributed Coordination: Drifting Environments and Constrained Communications.*, PhD Dissertation, Georgia Tech, 2014.
101. Y. Lim and J.S. Shamma, "Robustness of stochastic stability in game theoretic learning", *2013 American Control Conference*, June 2013.
102. M. Liu, Christopher Amato, Xuejun Liao, Jonathan P. How and Lawrence Carin. Stick-Breaking Policy Learning in Decentralized POMDPs. In *Proceedings of the Twenty-Fourth International Joint Conference on Artificial Intelligence (IJCAI-15)*, July 2015.
103. I. Lobel and A. Ozdaglar, "Distributed Subgradient Methods over Random Networks," *IEEE Transactions on Automatic Control* (to appear).
104. I. Lobel, A. Ozdaglar, and D. Feijer, "Distributed Multi-agent Optimization with State-Dependent Communication," LIDS report 2834, *Mathematical Programming* (to appear).
105. V.S. Mai and E.H. Abed, "Distributed Optimization Using Fast Consensus Algorithm, American Control Conference, Boston, July 6-8, 2016, submitted.
106. V.S. Mai and E.H. Abed, "Opinion dynamics with persistent leaders," 53rd IEEE Conference on Decision and Control, Los Angeles, December 15-17, 2014, pp. 2907-2913.
107. A. Mahajan, N. C. Martins, M. C. Rotkowitz and S. Yuksel, "Information Structures in Optimal Decentralized Control," *IEEE Conference on Decision and Control*, 2013.
108. H. Mahdavi, A. Beirami, B. Touri, and J.S. Shamma, Threshold policy for global games with noisy information sharing, IEEE Conference on Decision and Control, 2015.
109. V.S. Mai and E.H. Abed, "Opinion dynamics with persistent leaders", *IEEE Conference on Decision and Control*, 2014 (to appear).
110. A. Makhdoumi and A. Ozdaglar, "Broadcast-based Distributed Alternating Direction Method of Multipliers," *Proc. of Allerton*, 2014.
111. A. Makhdoumi and A. Ozdaglar, "Distributed Alternating Direction Method of Multipliers over Directed Networks," working paper, 2014.
112. J.R. Marden, S. Ruben, and L.Y. Pao, "Model-Free Approach to Wind Farm Control Using Game Theoretic Methods," *IEEE Transactions on Control Systems Technology*, special issue "to tame the wind: advanced control applications in wind energy," Volume 21, Number 4, pp. 1207-1214, 2013.
113. J.R. Marden and A. Wierman, "Distributed Welfare Games," *Operations Research*, Volume 61, Issue 1, pp. 155-168, 2013.
114. J.R. Marden and J.S. Shamma, Game theory and distributed control, in H.P. Young and S. Zamir (Eds): *Handbook of Game Theory*, Volume IV, Elsevier, 2015, pp. 861-899.
115. J.R. Marden and A. Wierman, "The Limitations of Utility Design for Multiagent Systems," *IEEE Transactions on Automatic Control*, Volume 58, Number 6, 2013.

116. J.R. Marden and J.S. Shamma, "Game Theory and Distributed Control," *Handbook of Game Theory, Volume IV*, 2013.
117. J. R. Marden, H. P. Young, and L. Pao, "Achieving Pareto Optimality through Distributed Learning", *SIAM Journal on Control and Optimization*. (to appear)
118. J. R. Marden, H. P. Young, and L. Pao, Achieving Pareto Optimality through Distributed Learning, *SIAM Journal on Control and Optimization* 52 (2014), 2753-2770.
119. I. Matei and J. S. Baras, "A Comparison Between Upper Bounds on Performance of two Consensus-based Distributed Optimization Algorithms", *Proceedings 3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys 2012)*, pp. 168-173, Santa Barbara, California, September 14-15, 2012.
120. I. Matei, C. Somarakis and J.S. Baras, "A Randomized Gossip Consensus Algorithm on Convex Metric Spaces", *Proceedings 51st IEEE Conference on Decision and Control, CDC2012*, 7425-7430, Maui, Hawaii, December 10-13, 2012.
121. I. Matei and J.S. Baras, "Distributed Algorithms for Optimization Problems with Equality Constraints," to appear in *Proceedings 2013 IEEE Conference on Decision and Control*, Florence, Italy, December 2013.
122. I. Matei and J.S. Baras, "The Asymptotic Consensus Problem on Convex Metric Spaces," to appear in *IEEE Transactions on Automatic Control*, 2013.
123. I. Matei, C. Somarakis and J.S. Baras, "A Generalized Gossip Algorithm on Convex Metric Spaces," submitted for journal publication, 2013.
124. I. Matei, C. Somarakis and J.S. Baras, "A Generalized Gossip Algorithm on Convex Metric Spaces," *IEEE Transactions on Automatic Control*, 2015. (to appear)
125. I. Matei and J.S. Baras, "Distributed Algorithms for Optimization Problems with Equality Constraints," *52nd IEEE Conference on Decision and Control*, pp. 2352-2357, Florence, Italy, December 10-13, 2013.
126. I. Matei and J.S. Baras, "The Asymptotic Consensus Problem on Convex Metric Spaces", *IEEE Transactions on Automatic Control*, 2015. (to appear)
127. I. Matei, J. S. Baras, M. Nabi and T. Kurtoglu, "An Extension of the Method of Multipliers for Distributed Nonlinear Programming", *53rd IEEE Conference on Decision Control (CDC 2014)*, Los Angeles, CA, 15-17 December 2014. (to appear)
128. I. Matei and J.S. Baras, "The Geometry of a Probabilistic Consensus of Opinion Algorithm," *Proceedings 2011 American Control Conference (ACC 2011)*, San Francisco, CA, June/July 2011, pp. 2198–2203.
129. I. Matei and J. S. Baras, "The Geometry of a Probabilistic Consensus of Opinion Algorithm," (submitted for journal publication).

130. I. Matei and J.S. Baras, "Power Allocation Policy for Distributed Estimation in Wireless Networks," *IEEE Conference on Decision and Control*, Orlando, Florida, December 2011, pp. 7275-7280.
131. I. Matei and J. S. Baras, "Convergence Results for the Agreement Problem on Markovian Random Topologies," *2011 IFAC World Conference*, Milano, Italy, August/September 2011, pp. 8860-8865.
132. I. Matei and J. S. Baras, "Performance Evaluation of the Consensus-based Distributed Sub-gradient Method Under Random Communication Topologies," *IEEE Journal of Selected Topics in Signal Processing*, Vol. 5, Issue 4, 2011, pp. 754-771.
133. I. Matei and J.S. Baras, "Consensus-Based Linear Distributed Filtering", *Automatica*, Vol. 48, June 2012, pp. 1776-1782.
134. I. Matei and J.S. Baras, "A Linear Distributed Filter Inspired by the Markovian Jump Linear System Filtering Problem," *Automatica*, Vol. 48, June 2012, pp. 1924-1928.
135. I. Matei and J.S. Baras, "The Asymptotic Consensus Problem on Convex Metric Spaces," (submitted for journal publication).
136. I. Matei, J.S. Baras and V. Srinivasan, "Trust-Based Multi-Agent Filtering for Increased Smart Grid Security," *Proceedings 20th Mediterranean Conference on Control and Automation*, Barcelona, Spain, July 2012, pp. 1266-1271,
137. I. Matei, C. Somarakis and J.S. Baras, "A Randomized Gossip Consensus Algorithm on Convex Metric Spaces", *IEEE Conference on Decision and Control*, Maui, Hawaii, December 2012 (to appear).
138. I. Matei and J. S. Baras, "A Performance Comparison Between Two Consensus-Based Distributed Optimization Algorithms," *3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys12)*, Santa Barbara, CA, September 2012 (to appear).
139. A. Menon and J. Baras, "Convergence Guarantees for a Decentralized Algorithm Achieving Pareto Optimality", *2013 American Control Conference*, pp. 1932-1937, Washington, DC, June 17 - 19, 2013.
140. A. Menon and J.S. Baras, "A Distributed Learning Algorithm with Bit-valued Communications for Multi-agent Welfare Optimization," *52nd IEEE Conference on Decision and Control*, pp. 2406-2411, Florence, Italy, December 10-13, 2013.
141. A. Menon and J. S. Baras, "Expander Families as Information Patterns for Distributed Control of Vehicle Platoons", *Proceedings 3rd IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys 2012)*, pp. 288-293, Santa Barbara, California, September 14-15, 2012.
142. A. Menon and J. Baras, "Convergence Guarantees for a Decentralized Algorithm Achieving Pareto Optimality", *Proceedings 2013 American Control Conference*, pp. 1932-1937, Washington, DC, June 17 - 19, 2013.

143. A. Menon and J. Baras, "A Distributed Learning Algorithm with Bit-valued Communications for Multi-agent Welfare Optimization," to appear in *Proceedings 2013 IEEE Conference on Decision and Control*, Florence, Italy, December 2013.
144. A. Menon and J. S. Baras, "Scheduling Laws for a Stochastic Distributed Learning Algorithm that Achieves Pareto Optimality", submitted for journal publication, 2013.
145. A. Menon and J. S. Baras, "A Distributed Learning Algorithm for Multi-agent Welfare Optimization," working paper, 2014.
146. A. Menon and J.S. Baras, "Collaborative Extremum Seeking for Welfare Optimization," *IEEE Conference on Decision and Control (CDC 2014)*, Los Angeles, California, December 15-17, 2014. (to appear)
147. H. H. Nax, M. Burton-Chellew, S. West, and H. P. Young, "Learning in a Black Box," working paper, 2014.
148. H. H. Nax and B.S.R. Pradelski, "Evolutionary dynamics and equitable core selection in assignment games," revise and resubmit, *International Journal of Game Theory*.
149. H. H. Nax, B. S.R. Pradelski, and H.P. Young, "Decentralized Dynamics to Optimal and Stable States in the Assignment Game," *IEEE Conference on Decision and Control*, 2013.
150. H. H. Nax, B. S.R. Pradelski, and H.P. Young, "Decentralized Dynamics and Equitable Core Selection in Assignment Games," submitted for journal publication, 2013.
151. H. H. Nax, B. S. R. Pradelski, and H. P. Young, "The evolution of core stability in decentralized matching markets", Discussion Paper 607, Department of Economics, University of Oxford.
152. H. H. Nax and B.S.R. Pradelski, Evolutionary dynamics and equitable core selection in assignment games, forthcoming in the *International Journal of Game Theory*, 2015.
153. F.A. Oliehoek, S. Whiteson, and Matthijs T. J. Spaan. "Exploiting Structure in Cooperative Bayesian Games", In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, August 2012.
154. F.A. Oliehoek, A.A. Gokhale, and Jie Zhang. "Reasoning about Advisors for Seller Selection in E-Marketplaces via POMDPs", In *15th International Workshop on Trust in Agent Societies*, pp. 67–78, June 2012.
155. F.A. Oliehoek and M.T.J. Spaan. "Tree-Based Solution Methods for Multiagent POMDPs with Delayed Communication", In *Proceedings of the National Conference on Artificial Intelligence*, July 2012.
156. F.A. Oliehoek, S. Witwicki, and L.P. Kaelbling. "Influence-Based Abstraction for Multiagent Systems", In *Proceedings of the National Conference on Artificial Intelligence*, July 2012.

157. F.A. Oliehoek and M.T.J. Spaan. “Tree-Based Solution Methods for Multiagent POMDPs with Delayed Communication”, In *AAMAS Workshop on Multi-agent Sequential Decision Making in Uncertain Domains*, pp. 48–55, June 2012. (longer version of AAAI)
158. F.A. Oliehoek and M.T.J. Spaan. “Tree-based Pruning for Multiagent POMDPs with Delayed Communication”, In *Proceedings of The International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp. 1229–1230, June 2012. Extended abstract.
159. F.A. Oliehoek. “Decentralized POMDPs”, In Wiering, Marco and van Otterlo, Martijn, editors, *Reinforcement Learning: State of the Art, Adaptation, Learning, and Optimization*, pp. 471–503, Springer Berlin Heidelberg, Berlin, Germany, 2012.
160. F.A. Oliehoek, S. Whiteson, and M.T.J. Spaan. “Exploiting Structure in Cooperative Bayesian Games”, In *Proceedings of the Conference on Uncertainty in Artificial Intelligence*, August 2012 (to appear).
161. F.A. Oliehoek, A.A. Gokhale, and J. Zhang. “Reasoning about Advisors for Seller Selection in E-Marketplaces via POMDPs”, In *15th International Workshop on Trust in Agent Societies*, pp. 67–78, June 2012.
162. F.A. Oliehoek and M.T.J. Spaan. “Tree-Based Solution Methods for Multiagent POMDPs with Delayed Communication”, In *Proceedings of the National Conference on Artificial Intelligence*, July 2012.
163. F.A. Oliehoek, S. Witwicki, and L.P. Kaelbling. “Influence-Based Abstraction for Multiagent Systems”, In *Proceedings of the National Conference on Artificial Intelligence*, July 2012.
164. F.A. Oliehoek and M.T.J. Spaan. “Tree-Based Solution Methods for Multiagent POMDPs with Delayed Communication”, In *AAMAS Workshop on Multi-agent Sequential Decision Making in Uncertain Domains*, pp. 48–55, June 2012. (longer version of AAAI)
165. F.A. Oliehoek and M.T.J. Spaan. “Tree-based Pruning for Multiagent POMDPs with Delayed Communication”, In *Proceedings of The International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp. 1229–1230, June 2012. Extended abstract.
166. F.A. Oliehoek. “Decentralized POMDPs”, In Wiering, Marco and van Otterlo, Martijn, editors, *Reinforcement Learning: State of the Art, Adaptation, Learning, and Optimization*, pp. 471–503, Springer Berlin Heidelberg, Berlin, Germany, 2012.
167. F.A. Oliehoek, Stefan Witwicki, and L.P. Kaelbling. “Heuristic Search of Multiagent Influence Space”, In *Proceedings of the European Workshop on Multi-agent Systems*, 2011.
168. F. A. Oliehoek and C. Amato, “Best-Response Bayesian Reinforcement Learning for Multiagent Systems with State Uncertainty”, *Proceedings of the Workshop on Multi-Agent Sequential Decision Making in Uncertain Domains (MSDM)* at the Thirteenth International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS-14), May 2014.

169. S. Omidshafiei, Ali-akbar Agha-mohammadi, Christopher Amato and Jonathan P. How. Decentralized Control of Partially Observable Markov Decision Processes using Belief Space Macro-actions. In *Proceedings of the 2015 IEEE International Conference on Robotics and Automation (ICRA-15)*, May 2015.
170. B.S.R. Pradelski, Decentralized Dynamics and Fast Convergence in the Assignment Game, 16th ACM Conference on Economics and Computation (EC), 2015.
171. B.S.R. Pradelski. The Dynamics of Social Influence, University of Oxford Department of Economics Discussion Paper 742.
172. S. Park and N.C. Martins, "Necessary and sufficient conditions for the stabilizability of a class of LTI distributed observers", *IEEE Conference on Decision and Control*, Maui, HI, December 2012 (to appear).
173. S. Park and N. C. Martins, "A Class of LTI Distributed Observers for LTI Plants: Necessary and Sufficient Conditions for Asymptotic State Omniscience," working paper, 2014.
174. S. Park, *Distributed Estimation And Stability Of Evolutionary Games For Large Populations, With Applications To Cyber Physical Systems*, Ph.D. Thesis proposal, ECE Department, University of Maryland at College Park, June 2014.
175. S. Park and N. C. Martins, "Individually Optimal Solutions to a Remote State Estimation Problem with Communication," *IEEE Conference on Decision and Control*, 2014. (to appear)
176. S. Park and N. C. Martins, "Necessary and Sufficient Conditions for the Stabilizability of a Class of LTI Distributed Observers," *Proceedings of the IEEE Conference on Decision and Control*, 2012.
177. S. Park and N. C. Martins, Individually optimal solutions to a remote state estimation problem with communication costs, *Proceedings of the IEEE CDC*, pp. 4014- 4019, 2014 (under review) A draft is available at arXiv:1401.0926.
178. S. Park and N. C. Martins, Design of Distributed LTI Observers for State Omniscience, *IEEE Transactions on Automatic Control*.
179. S. Park, *Distributed Estimation and Stability of Evolutionary Games for Large Populations*, Ph.D. Thesis Proposal, ECE Department, UMCP, June 2014 (Defense is scheduled for Nov 6, 2015)
180. D. Pickem, M. Egerstedt, and J.S. Shamma, A game-theoretic formulation of the homogeneous self-reconfiguration problem, *IEEE Conference on Decision and Control*, 2015.
181. D. Pickem, M. Egerstedt, and J.S. Shamma, "Complete heterogeneous self-reconfiguration: Deadlock avoidance using hole-free assemblies", *4th IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys'13)*, Koblenz, Germany, September 2013.

182. D. Pickem, M. Egerstedt, and J.S. Shamma, "Complete heterogeneous self-reconfiguration: Deadlock avoidance using hole-free assemblies", *4th IFAC Workshop on Distributed Estimation and Control in Networked Systems (NecSys'13)*, September 2013.
183. G. Piliouras and J.S. Shamma, "Optimization despite chaos: Convex relaxations to complex limit sets via Poincaré recurrence", *ACM-SIAM Symposium on Discrete Algorithms (SODA14)*, Portland, OR, January 2014.
184. G. Piliouras, C. Nieto-Granda, H. Christensen, J.S. Shamma, "Persistent patterns: Multi-agent learning beyond equilibrium and utility", *13th International Conference on Autonomous Agents and Multiagent Systems (AAMAS2014)*, Paris, France, May 2014.
185. G. Piliouras, E. Nikolova, and J.S. Shamma, "Risk sensitivity of price of anarchy under uncertainty", *14th ACM Conference on Electronic Commerce (EC'13)*, June 2013.
186. B.S.R. Pradelski, "Evolutionary Dynamics and Fast Convergence in the Assignment Game", working paper, 2014.
187. S. V. Rakovic, B. Marinkovic and J. S. Baras, "Safety, Invariance and Nagumo Theorem for a Class of Constrained Nonlinear Discrete-Time Dynamics," working paper, 2014.
188. S.V. Rakovic, I. Matei and J.S. Baras, "Reachability Analysis for Linear Discrete Time Set-Dynamics Driven by Random Convex Compact Sets," *Proceedings 51st IEEE Conference on Decision and Control*, CDC2012, pp. 4751-4756, Maui, Hawaii, December 10-13, 2012.
189. S.V. Rakovic, M. Lazar, B. Marinkovic and J.S. Baras, "Safety, Invariance and Nagumo Theorem for a Class of Constrained Nonlinear Discrete Time Dynamics," submitted for journal publication, 2013.
190. S. V. Rakovic, I. Matei and J.S. Baras, "Reachability Analysis for Linear Discrete Time Set-Dynamics Driven by Random Convex Compact Sets," *IEEE Conference on Decision and Control*, Maui, Hawaii, December 2012.
191. M.C. Rotkowitz and N.C. Martins, "On the nearest quadratically invariant information constraint", *IEEE Transactions on Automatic Control*, 57(5), pp. 1314-1319.
192. M. C. Rotkowitz and N. C. Martins, "On the Nearest Quadratically Invariant Information Constraint," *IEEE Transactions on Automatic Control*, Vol 57, No 5, 2012, pp. 1314-1319.
193. S. Sabau and N. C. Martins, "Stabilizability and Norm-Optimal Control Design subject to Sparsity Constraints," to appear in *IEEE Transactions on Automatic Control*, 2013.
194. S. Sabau and N. C. Martins, "Youla-like Parametrizations Subject to QI Subspace Constraints," *IEEE Transactions on Automatic Control*, Vol. 59, No. 6, 2014, pp. 1411–1423.
195. S. Sabau, and N.C. Martins, "Necessary and sufficient conditions for stabilizability subject to quadratic invariance", *IEEE Conference on Decision and Control*, Orlando, FL, December 2011.

196. S. Sabau and N.C. Martins, "Necessary and sufficient conditions for stabilizability subject to sparsity constraints", (submitted for journal publication).
197. S. Sabau and N.C. Martins, "Optimal performance improvement of diagonal controllers with fixed entries", *The Fifteenth Yale Workshop on Adaptive and Learning Systems* New Haven, CT, June 2011.
198. S. Sabau, *Synthesis of Sparsity Constrained Controllers: New Methods and Algorithm*, Ph.D. Thesis, U. Maryland at College Park, 2011.
199. S. Sabau and N. C. Martins, Youla-like parametrizations subject to QI Subspace Constraints, *IEEE Transactions on Automatic Control*, Vol. 59, No. 6, pp. 1411-1423, 2014
200. K. Savla, G. Como, and M.A. Dahleh, "Robust Oblivious Routing for Network Flows under Cascading Failures", *IEEE Conference on Decision and Control*, Los Angeles, CA, December 2014. (to appear)
201. S. Shtern, E. Wei, and A. Ozdaglar, "Distributed Alternating Direction Method of Multipliers: Performance and Network Effects," working paper, 2014.
202. C. Somarakis and J. S. Baras, "The Consensus Problem Under Vanishing Communication Weights", *Proceedings of the 20th International Symposium on Mathematical Theory of Networks and Systems (MTNS 2012)*, Melbourne, Australia, July 2012.
203. E. Stai, J. S. Baras and S. Papavassiliou, "Throughput-Delay Trade-off in Wireless Multi-hop Networks via Greedy Hyperbolic Embedding", *20th International Symposium on Mathematical Theory of Networks and Systems (MTNS 2012)*, Melbourne, Australia, 2012.
204. E. Stai, J.S. Baras and S. Papavassiliou, "A Class of Backpressure Algorithms for Networks Embedded in Hyperbolic Space with Controllable Delay-Throughput Trade-off," *15th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2012)*, Paphos, Cyprus, October 2012 (to appear).
205. C. Somarakis and J.S. Baras, "LTI Continuous Time Consensus Dynamics with Multiple and Distributed Delays," *21st Mediterranean Conference on Control and Automation (MED)*, pp. 1495-1501, Platanias-Chania, Greece, June 25–28, 2013.
206. C. Somarakis and J.S. Baras, "On the Connectivity Assumption of Non-Linear Flocking Models," *18th International Conference of Methods and Models in Automation and Robotics (MMAR 2013)*, pp. 800-805, Midzysdroje, Poland, August 26-29, 2013.
207. C. Somarakis and J.S. Baras, "On The Convergence Rate of Non-Linear Consensus Networks with Delays," *9th IFAC Symposium on Nonlinear Control Systems (NOLCOS 2013)*, pp. 705-710, Toulouse, France, August 2013.
208. C. Somarakis and J.S. Baras, "Fixed Point Theory Approach to Exponential Convergence In LTV Continuous Time Consensus Dynamics with Delays," *2013 SIAM Conference on Control and Its Applications*, pp. 129-136, San Diego, California, July 8-10, 2013.

209. C. Somarakis and J.S. Baras, "Asymptotic Consensus Solutions in Non-Linear Delayed Distributed Algorithms Under Deterministic & Stochastic Perturbations," *52nd IEEE Conference on Decision and Control*, pp. 7821-7826, Florence, Italy, December 10-13, 2013.
210. C. Somarakis, J.S. Baras and E. Paraskevas, "On the Dynamics of $\dot{x} = -ax + ax_t$ ", working paper, 2014.
211. C. Somarakis, J.S. Baras and E. Paraskevas, "Stability by Fixed Point Theory in Time Delayed Distributed Consensus Networks," working paper, 2014.
212. W.H. Sandholm and M. Staudigl, Large Deviations and Stochastic Stability in the Small Noise Double Limit, forthcoming in *Theoretical Economic*, 2015.
213. E. Stai, S. Papavassiliou and J.S. Baras, "Performance-Aware Cross-Layer Design in Wireless Multihop Networks via a Weighted Backpressure Approach," *IEEE/ACM Transactions on Networking*, 2015. (to appear)
214. E. Stai, J.S. Baras and S. Papavassiliou, "Social Networks over Wireless Networks", *IEEE Conference on Decision and Control*, Maui, Hawaii, December 2012.
215. E. Stai, J.S. Baras and S. Papavassiliou, "A Class of Backpressure Algorithms for Networks Embedded in Hyperbolic Space with Controllable Delay-Throughput Trade-off," *Proceedings 15th ACM International Conference on Modeling, Analysis and Simulation of Wireless and Mobile Systems (MSWiM 2012)*, Paphos, Cyprus, October 21-25, 2012.
216. E. Stai, J.S. Baras and S. Papavassiliou, "Social Networks over Wireless Networks," *Proceedings 51st IEEE Conference on Decision and Control, CDC2012*, pp. 2696-2703, Maui, Hawaii, December 10-13, 2012.
217. E. Stai, S. Papavassiliou and J.S. Baras, "Performance-Aware Cross-Layer Design in Wireless Multihop Networks via a Weighted Backpressure Approach," submitted for journal publication, 2013.
218. M. Staudigl, "A Limit Theorem for Markov Decision Processes," submitted to *Mathematics of Operations Research*, 2012.
219. B. Touri and J.S. Shamma, "Global games with noisy sharing of information", *53rd IEEE Conference on Decision and Control*, Los Angeles, CA, December 2014. (to appear)
220. P. Tschirhart, A. Sussman and E.H. Abed, "Using participation factors to improve the consistency and accuracy of Prony analysis for voltage stability monitoring applications," Paper No. ISGT2014-000055, *Proc. IEEE PES Conference on Innovative Smart Grid Technologies (ISGT)*, Washington, DC, Feb. 19-22, 2014.
221. P.K. Tschirhart, *Implications of the Dichotomy of Modal Participation Factors for Monitoring and Control of Electric Power Networks*, M.S. Thesis, University of Maryland, College Park, 2013.

222. M. M. Vasconcelos, and N. C. Martins, “Estimation Over the Collision Channel: Structural Results,” *Allerton Conference on Communication, Control and Computing*, 2013.
223. M. Vasconcelos and N. C. Martins, Estimation over the Collision Channel: Structural Results, *Proceedings of the Annual Allerton Conference*, pp. 1114–1119, 2013.
224. K. Voevodski, M.F. Balcan, H. Röglin, S.-H. Teng, and Y. Xia, “Active clustering of biological sequences”, *Journal of Machine Learning Research*, 2012, pp. 203–225.2
225. D. Ward, N. C. Martins and B. M. Sadler, Optimal Remote Estimation over Action Dependent Switching Channels: Managing Workload and Bias of a Human Operator, *Proceedings of the ACC*, pp. 3168–3174, 2015.
226. E. Wei and A. Ozdaglar, “On the Convergence of Asynchronous Distributed Alternating Direction Method of Multipliers,” working paper, 2013.
227. X. Wei and E.H. Abed, “Analysis of epidemics in a network of interacting communities,” working paper, 2014.
228. E. Wei and A. Ozdaglar, “On the $O(1/k)$ Convergence of Asynchronous Distributed Alternating Direction Method of Multipliers, submitted for journal publication, 2013.
229. E. Wei and A. Ozdaglar, “Distributed Alternating Direction Method of Multipliers,” *Proc. of IEEE Conference on Decision and Control*, 2012.
230. X. Wei, Diffusion Dynamics in Networked Communities, Ph.D. Dissertation, Dec. 2015, University of Maryland, College Park (in preparation; Advisor: E.H. Abed).
231. X. Wei and E.H. Abed, “Analysis of epidemics in a network of interacting communities,” in preparation.
232. X. Wei and E.H. Abed, “Maximization of epidemic spread in networks of interacting communities,” in preparation.
233. X. Wei and E.H. Abed, “Percolation theory analysis of epidemic spread in networks of interacting communities,” in preparation.
234. S. Witwicki, F.A. Oliehoek, and L.P. Kaelbling. “Heuristic Search of Multiagent Influence Space”, In *Proceedings of The International Joint Conference on Autonomous Agents and Multi Agent Systems*, pp. 973–981, June 2012
235. H.P. Young, “The dynamics of social innovation”, *Proceedings of the National Academy of Sciences*, 2011, vol. 108 (4), pp. 21285–21291.
236. Y. Zhou, J. Moschler and J. S. Baras, “A Systems Engineering Approach to Collaborative Coordination of UAVs in the NAS with Safety Guarantees,” *2014 Integrated Communications Navigation and Surveillance (ICNS) Conference*, Herndon, Virginia, April 8-10, 2014.

7 Honors & Awards

- Eyad Abed
 - Appointed as a DeTao Master in Systems and Control: New Information Technology, Beijing DeTao Masters Academy, a high-level, multi-disciplined, application-oriented higher education institution in Shanghai, China (January 2014).
- Maria-Florina Balcan
 - Sloan Research Fellowship, 2014
 - Program Chair for the 27th Annual Conference on Learning Theory (COLT) 2014.
- John S. Baras
 - Elected Fellow of the Society for Industrial and Applied Mathematics (SIAM) “for contributions to systems theory, stochastic control, and communication networks”
 - Tage Erlander Guest Professorship by the Swedish Research Council
 - Hans Fischer Senior Fellowship from the Institute of Advanced Studies (IAS) of the Technical University of Munich (TUM)
 - Best Student Paper Finalist, 2013 *American Control Conference*
 - Best Student Paper Finalist, 2013 *IEEE Conference on Decision and Control*
 - Dissertation Fellowship for advisee Anup Menon, U of Maryland College Park.
 - Plenary speaker: General Electric Global Research Controls Symposium, 9th Carnegie Mellon University Electricity Conference, 2nd Hellenic Forum for Science, Innovation and Technology.
 - Invited speaker: University of Pennsylvania, United Technologies Research Center, Xerox PARC, UC Berkeley TRUST Center, Stanford University, University of Notre Dame, NIST, University of Minnesota, University of Thessaly, University of Thessaloniki.
- Munther Dahleh
 - William Coolidge Institute Chair, MIT
 - Plenary speaker: Mediterranean Conference on Control and Automation, June, 2014.
 - Invited speaker: Workshop on “Robustness in Identification, Control ... and Beyond”, Siena, Italy, September 2014.
- Leslie Kaelbling
 - Best-paper award at the 2014 International Conference on Autonomous Agents and Multi-Agent Systems (AAMAS).
 - Nominated for the best paper award at the 2015 Robotics: Science and Systems Conference (RSS).

- Nuno Martins
 - Program vice-chair for the IEEE CDC in 2013 and 2014
 - Director of the University of Maryland's Robotics Center 2012-2014
 - Co-organizer (jointly with Anders Rantzer and Naomi Leonard) of the IMA workshop on Distributed Control and Decision Making Over Networks, September 28 – October 2, 2015
- Asu Ozdaglar
 - Recipient of Spira teaching award from MIT EECS
 - Plenary speaker: *Modeling and Optimization: Theory and Applications (MOPTA)*, Lehigh University.
 - Invited talk: “Stochastic gradient methods”, IPAM, UCLA.
- Jeff Shamma:
 - *Plenary Speaker*: IEEE Conference on Decision and Control, December 2013.
 - *Invited Speaker*: KAUST, February 2014; Cornell University, April 2014; University of Montreal, May 2014; Workshop on “Robustness in Identification, Control ... and Beyond”, Siena, Italy, September 2014.
- Peyton Young
 - Plenary speaker: International Game Theory Conference, State University of New York, Stony Brook, July 2013; Conference on Mathematics of Social Learning, UCLA, January 2014; Sackler Colloquium of the National Academy of Sciences, January 2014; Conference on Financial Interconnectedness, International Monetary Fund, May 2014; 70 Years of Game Theory, Tinbergen Institute, Amsterdam, June 2014.
 - Invited speaker: American Economic Association Annual Meeting, January 2014; University of Glasgow, February 2014; University of Cambridge, May 2014.
 - Plenary Speaker: International Game Theory Conference, State University of New York at Stony Brook, July 2014
 - Invited speaker: University of Cambridge, August 2014; Stanford University, September 2014; University of Rotterdam, May 2015.

8 Transitions

- Christopher Amato became faculty at the University of New Hampshire in January, 2015. Peyton Young: Collaboration with control theorists Jason Marden and Lucy Pao (University of Colorado).
- Discovery of Fast network consensus algorithm and application to fast distributed optimization and consensus prediction from an observed subset of system nodes.

- Discovery of Analytical conditions for epidemic outbreak for networks of interacting communities, and validation using percolation theory for large networks and simulation.
- Ph.D. student Xiaoya Wei (University of Maryland) worked as an intern at Facebook in summer 2015; he is considering joining the company upon completion of his Ph.D. (expected Dec. 2015).

9 New Discoveries

None to report.

1.

1. Report Type

Final Report

Primary Contact E-mail

Contact email if there is a problem with the report.

feron@gatech.edu

Primary Contact Phone Number

Contact phone number if there is a problem with the report

617-413-7666

Organization / Institution name

Georgia Tech Research Corporation

Grant/Contract Title

The full title of the funded effort.

DISTRIBUTED LEARNING AND INFORMATION DYNAMICS
IN NETWORKED AUTONOMOUS SYSTEMS

Grant/Contract Number

AFOSR assigned control number. It must begin with "FA9550" or "F49620" or "FA2386".

FA9550-09-1-0538

Principal Investigator Name

The full name of the principal investigator on the grant or contract.

Eric M Feron

Program Manager

The AFOSR Program Manager currently assigned to the award

Fariba Fahroo

Reporting Period Start Date

07/01/2009

Reporting Period End Date

06/30/2015

Abstract

This project is motivated by the need to enable advanced operations of teams of autonomous vehicles to learn and adapt to uncertain and hostile environments under effective utilization of communications resources. Of particular interest is the interplay between distributed learning and information dynamics. Distributed learning refers to a collection of interacting agents with limited local processing, information, and communications, all seeking to achieve a global objective in an uncertain and possibly hostile environment. Information dynamics refers to the architecture, either inherited or designed, of information flow among the distributed agents. The interplay of distributed learning algorithms and information dynamics can have dramatic effects on the efficiency of the collective. Specific research thrusts include: i) online resource allocation, ii) networked operations, and iii) evolving and uncertain operations.

Distribution Statement

This is block 12 on the SF298 form.

Distribution A - Approved for Public Release

Explanation for Distribution Statement

If this is not approved for public release, please provide a short explanation. E.g., contains proprietary information.

SF298 Form

Please attach your [SF298](#) form. A blank SF298 can be found [here](#). Please do not password protect or secure the PDF. The maximum file size for an SF298 is 50MB.

[Final_report_cover_page.pdf](#)

Upload the Report Document. File must be a PDF. Please do not password protect or secure the PDF. The maximum file size for the Report Document is 50MB.

[AFOSR Grant FA95500910538.pdf](#)

Upload a Report Document, if any. The maximum file size for the Report Document is 50MB.

Archival Publications (published) during reporting period:

See report

Changes in research objectives (if any):

N/A

Change in AFOSR Program Manager, if any:

N/A

Extensions granted or milestones slipped, if any:

Grant expiration extended from June 30, 2014 to June 30, 2015.

AFOSR LRIR Number

LRIR Title

Reporting Period

Laboratory Task Manager

Program Officer

Research Objectives

Technical Summary

Funding Summary by Cost Category (by FY, \$K)

	Starting FY	FY+1	FY+2
Salary			
Equipment/Facilities			
Supplies			
Total			

Report Document

Report Document - Text Analysis

Report Document - Text Analysis

Appendix Documents

2. Thank You

E-mail user

Nov 09, 2015 12:40:35 Success: Email Sent to: feron@gatech.edu